

US007565016B2

(12) **United States Patent**  
**Hua et al.**

(10) **Patent No.:** **US 7,565,016 B2**  
(45) **Date of Patent:** **Jul. 21, 2009**

(54) **LEARNING-BASED AUTOMATIC  
COMMERCIAL CONTENT DETECTION**

(75) Inventors: **Xian-Sheng Hua**, Beijing (CN); **Lie Lu**,  
Beijing (CN); **Mingjing Li**, Beijing  
(CN); **Hong-Jiang Zhang**, Beijing (CN)

(73) Assignee: **Microsoft Corporation**, Redmond, WA  
(US)

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 0 days.

(21) Appl. No.: **11/623,304**

(22) Filed: **Jan. 15, 2007**

(65) **Prior Publication Data**

US 2007/0112583 A1 May 17, 2007

**Related U.S. Application Data**

(63) Continuation of application No. 10/368,235, filed on  
Feb. 18, 2003, now Pat. No. 7,164,798.

(51) **Int. Cl.**

**G06K 9/72** (2006.01)

**H04H 9/00** (2006.01)

(52) **U.S. Cl.** ..... **382/229; 725/22**

(58) **Field of Classification Search** ..... **382/229;**  
**725/22**

See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,333,091 A 7/1994 Iggulden et al.  
5,442,633 A 8/1995 Perkins et al.  
5,497,430 A 3/1996 Sadovnik et al.  
5,530,963 A 6/1996 Moore et al.  
5,625,877 A 4/1997 Dunn et al.

5,642,294 A 6/1997 Taniguchi et al.  
5,659,685 A 8/1997 Williams et al.  
5,710,560 A 1/1998 Cohn  
5,745,190 A 4/1998 Ioka  
5,751,378 A 5/1998 Chen et al.  
5,774,593 A 6/1998 Zick et al.  
5,778,137 A 7/1998 Nielsen et al.  
5,801,765 A 9/1998 Gotoh et al.

(Continued)

**FOREIGN PATENT DOCUMENTS**

EP 0597450 5/1994

(Continued)

**OTHER PUBLICATIONS**

U.S. Appl. No. 10/285,933, filed Nov. 1, 2002 inventors Yu-Fei Ma et  
al. entitled "Systems and Methods for Generating a Motion Attention  
Mode".

(Continued)

*Primary Examiner*—Brian P Werner

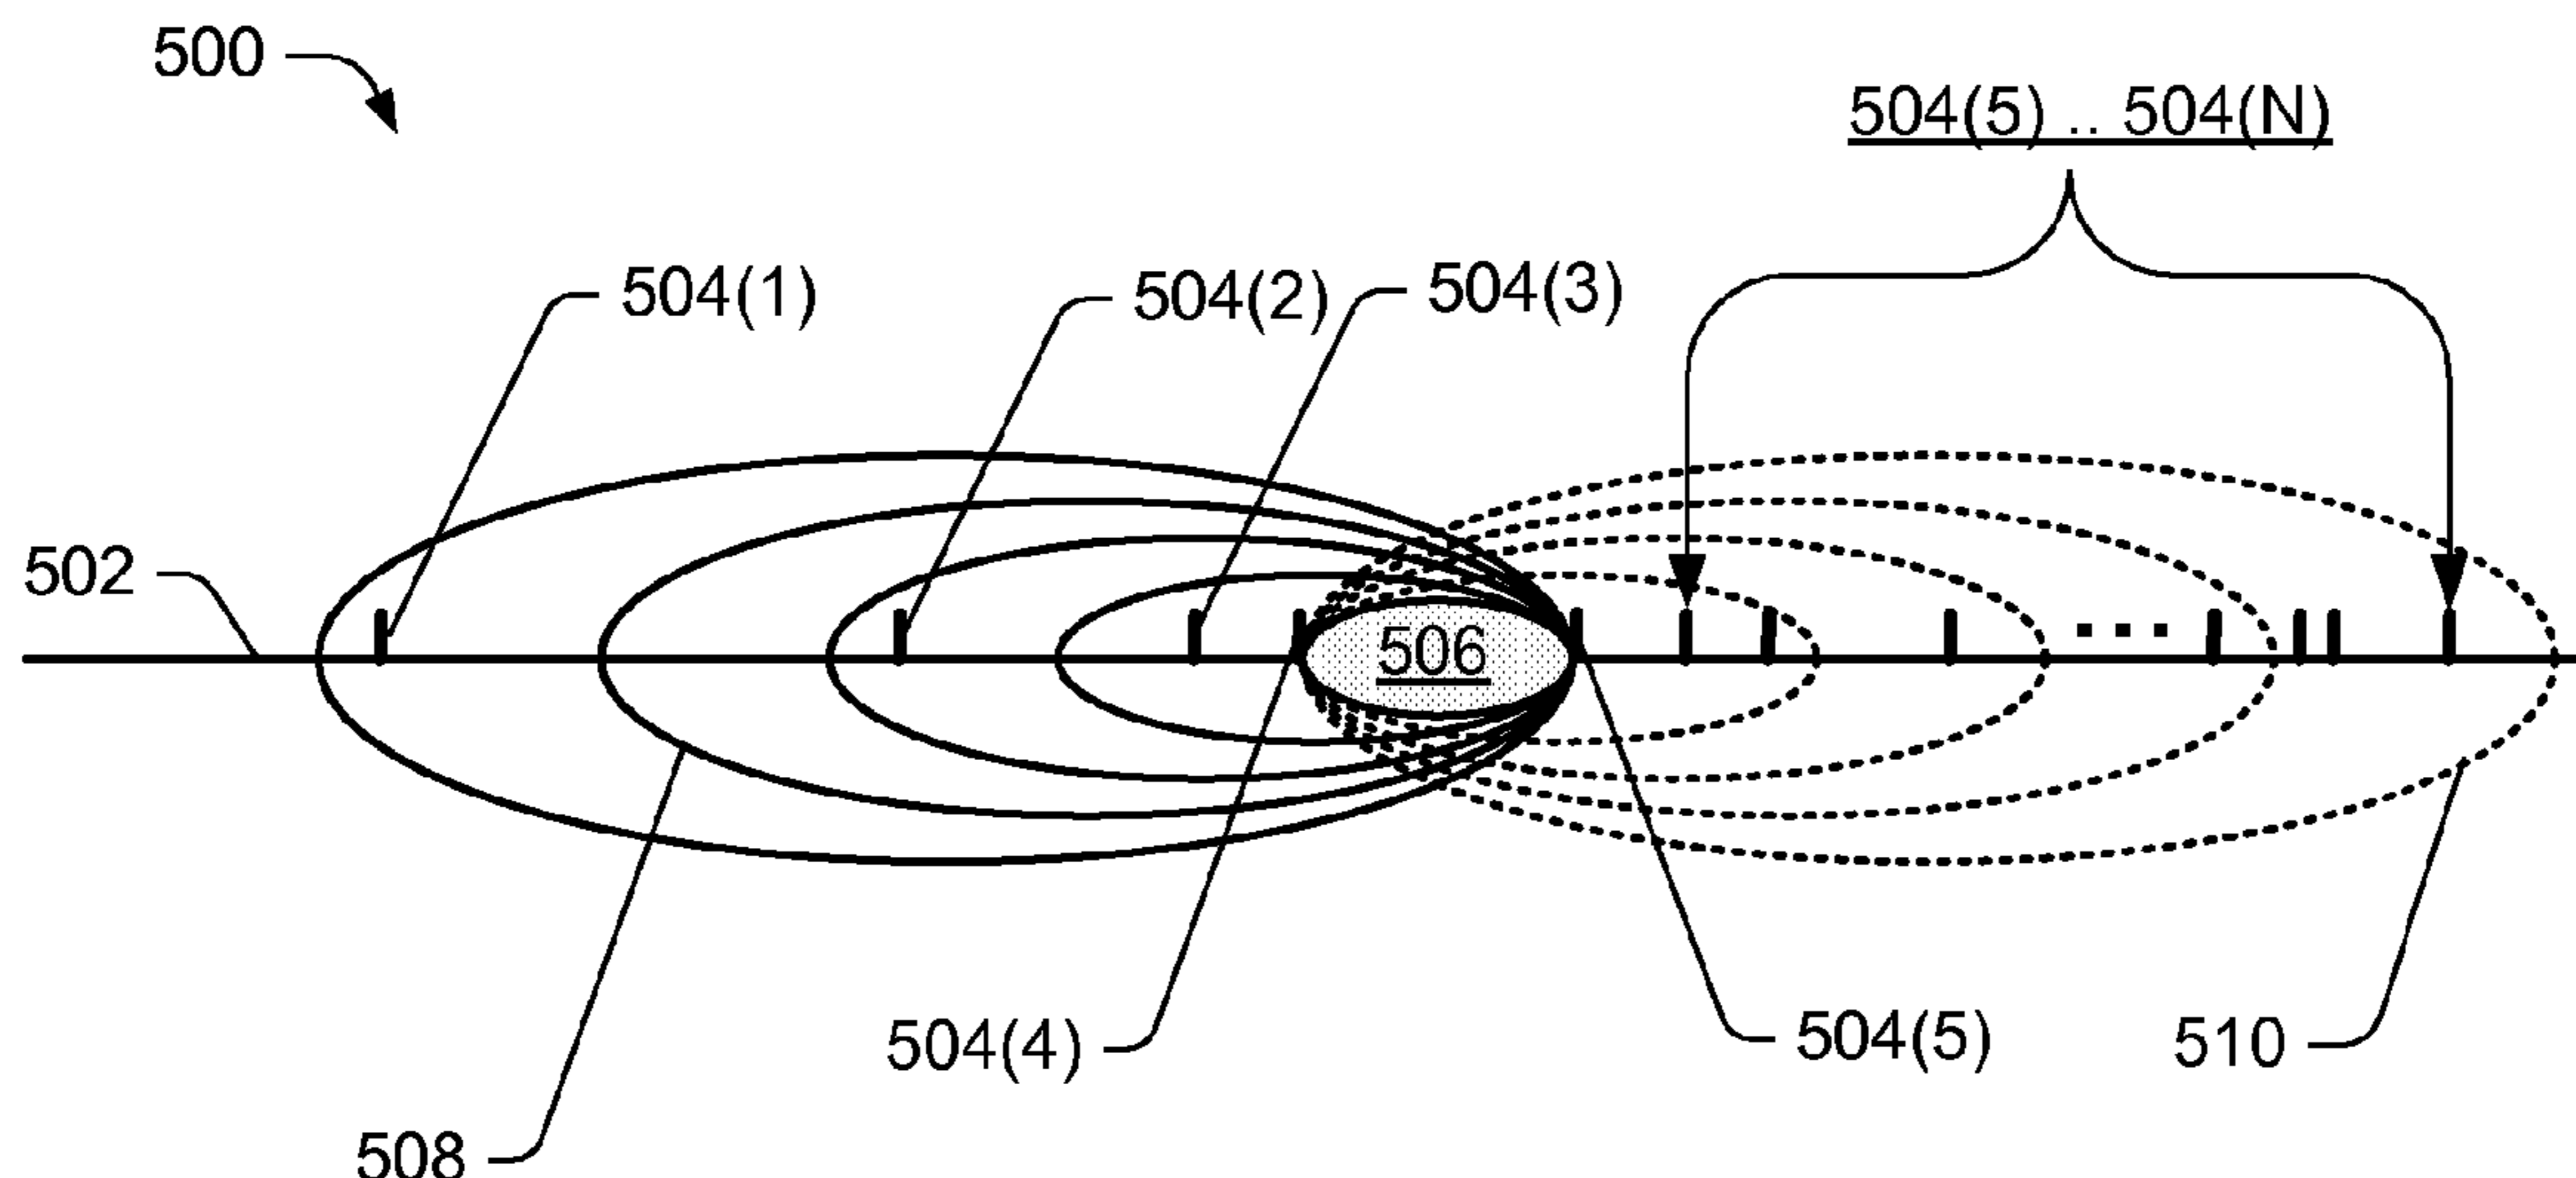
*Assistant Examiner*—Kathleen S Yuan

(74) *Attorney, Agent, or Firm*—Lee & Hayes, PLLC

(57) **ABSTRACT**

Systems and methods for learning-based automatic commer-  
cial content detection are described. In one aspect, the sys-  
tems and methods include a training component and an analy-  
zing component. The training component trains a  
commercial content classification model using a kernel sup-  
port vector machine. The analyzing component analyzes pro-  
gram data such as video and audio data using the commercial  
content classification model and one or more of single-side  
left neighborhood(s) and right neighborhood(s) of program  
data segments. Based on this analysis, each of the program  
data segments are classified as being commercial or non-  
commercial segments.

**10 Claims, 6 Drawing Sheets**



U.S. PATENT DOCUMENTS

5,835,163 A 11/1998 Liou et al.  
 5,884,056 A 3/1999 Steele  
 5,900,919 A 5/1999 Chen et al.  
 5,901,245 A 5/1999 Warnick et al.  
 5,911,008 A 6/1999 Niikura et al.  
 5,920,360 A 7/1999 Coleman, Jr.  
 5,952,993 A 9/1999 Matsuda et al.  
 5,956,026 A 9/1999 Ratakonda  
 5,959,697 A 9/1999 Coleman, Jr.  
 5,966,126 A 10/1999 Szabo  
 5,983,273 A 11/1999 White et al.  
 5,990,980 A 11/1999 Golin  
 5,995,095 A 11/1999 Ratakonda  
 6,020,901 A 2/2000 Lavelle et al.  
 6,047,085 A 4/2000 Sato et al.  
 6,100,941 A \* 8/2000 Dimitrova et al. .... 348/700  
 6,166,735 A 12/2000 Dom et al.  
 6,168,273 B1 1/2001 Dupraz et al.  
 6,182,133 B1 1/2001 Horvitz  
 6,232,974 B1 5/2001 Horvitz et al.  
 6,236,395 B1 5/2001 Sezan et al.  
 6,282,317 B1 8/2001 Luo et al.  
 6,292,589 B1 9/2001 Chow et al.  
 6,307,550 B1 10/2001 Chen et al.  
 6,353,824 B1 3/2002 Boguraev et al.  
 6,408,128 B1 6/2002 Abecassis  
 6,421,675 B1 7/2002 Ryan et al.  
 6,462,754 B1 10/2002 Chakraborty et al.  
 6,466,702 B1 10/2002 Atkins et al.  
 6,473,778 B1 10/2002 Gibbon  
 6,581,096 B1 6/2003 Cottrille et al.  
 6,616,700 B1 9/2003 Thum et al.  
 6,622,134 B1 9/2003 Sorkin  
 6,643,643 B1 11/2003 Lee et al.  
 6,643,665 B2 11/2003 Kimbell et al.  
 6,658,059 B1 12/2003 Iu et al.  
 6,661,468 B2 12/2003 Alten et al.  
 6,670,963 B2 12/2003 Osberger  
 6,714,909 B1 3/2004 Gibbon et al.  
 6,773,778 B2 8/2004 Onozawa et al.  
 6,792,144 B1 9/2004 Yan et al.  
 6,807,361 B1 10/2004 Girgensohn et al.  
 6,870,956 B2 3/2005 Qi et al.  
 6,934,415 B2 8/2005 Stentiford  
 7,006,091 B2 2/2006 Masera et al.  
 7,055,166 B1 5/2006 Logan et al.  
 7,062,705 B1 6/2006 Kirkwood et al.  
 7,065,707 B2 6/2006 Chen et al.  
 7,116,716 B2 10/2006 Ma et al.  
 7,248,294 B2 7/2007 Slatter  
 2001/0023450 A1 9/2001 Chu  
 2001/0047355 A1 11/2001 Anwar  
 2002/0069218 A1 6/2002 Sull et al.  
 2002/0100052 A1 7/2002 Daniels  
 2002/0157116 A1 10/2002 Jasinski  
 2002/0166123 A1 11/2002 Schrader et al.  
 2003/0033347 A1 \* 2/2003 Bolle et al. .... 709/107  
 2003/0123850 A1 7/2003 Jun et al.  
 2003/0152363 A1 8/2003 Jeannin et al.  
 2003/0210886 A1 11/2003 Li et al.  
 2003/0237053 A1 12/2003 Chen et al.  
 2004/0040041 A1 2/2004 Crawford  
 2004/0068481 A1 4/2004 Seshadri et al.  
 2004/0078357 A1 4/2004 LaChapelle et al.  
 2004/0078382 A1 4/2004 Mercer et al.  
 2004/0078383 A1 4/2004 Mercer et al.  
 2004/0085341 A1 5/2004 Hua et al.  
 2004/0088726 A1 5/2004 Ma et al.  
 2004/0165784 A1 8/2004 Xie et al.  
 2004/0184776 A1 9/2004 Inoue et al.  
 2006/0239644 A1 10/2006 Barbieri

2007/0027754 A1 2/2007 Collins et al.  
 2007/0060099 A1 3/2007 Ramer et al.

FOREIGN PATENT DOCUMENTS

EP 1168840 1/2002  
 EP 1213915 6/2002  
 GB 2356080 5/2001  
 JP 2002238027 8/2002  
 KR 20020009089 A 2/2002  
 KR 20040042449 A 5/2004  
 WO WO0028467 A1 5/2000

OTHER PUBLICATIONS

U.S. Appl. No. 10/286,053, filed Nov. 1, 2002 inventors Yu-Fei Ma et al. entitled "Systems and Methods for Generating a Comprehensive User Attention Model".  
 U.S. Appl. No. 10/286,348, filed Nov. 1, 2002 inventors Xian-Sheng Hua et al. entitled "Systems and Methods for Automatically Editing a Video".  
 U.S. Appl. No. 10/371,125, filed Feb. 20, 2003 Inventors Wei-Ying Ma et al. Entitled "Systems and Methods for Enhanced Image Adaptation".  
 "Advanced Digital Video Storage and On-line Retrieval System", Advisor Project, Jun. 30, 2002, retrieved from the Internet: [http://advisor.matrasi-tis.fr/DUP\\_workshop\\_sheet.pdf](http://advisor.matrasi-tis.fr/DUP_workshop_sheet.pdf).  
 Ahmad, "VISIT: A Neural Model of Covert Visual Attention" Advances in Neural Information Processing Systems v. 4 1991 pp. 420-427.  
 Baluja et al., "Expectation-Based Selective Attention for Visual Monitoring and Control of a Robot Vehicle" Robotics and Autonomous System v. 22 No. 3-4 Dec. 1997 pp. 329-344.  
 Bertini, et al., "Indexing for Reuse of TV News Shots" Pattern Recognition, Elsevier, Kidlington, GB, vol. 35, No. 3, Mar. 2002, pp. 581-591.  
 Bertini et al., "Indexing for Resue of TV News Shots" Pattern Recognition, Elsevier, Kidlington, GB, vol. 35, No. 3, Mar. 2002, pp. 581-591.  
 Chen et al., "A Visual Attention Model for Adapting Images on Small Displays" ACM Multimedia Systems Journal 2003.  
 Chen et al., "Function-based Object Model Towards Website Adaptation" (2001) Proc. of the 10th Int. wWW Conf. pp. 1-21.  
 Cherry, "Pesky Home Networks Trouble cable Behemoths" IEEE Spectrum Apr. 2002 pp. 60-61.  
 Christopoulos et al., "The JPEG2000 Still Image Coding System: An Overview" IEEE Transactions on Consumer Electronics vol. 46 No. 4 pp. 1103-1127 Nov. 2000.  
 Deng et al., "Peer Group Filtering and Perceptual Color Image Quantization" Proceedings of IEEE International Symposium on Circuits and systems vol. 4 1999 pp. 21-24.  
 Divakaran et al., "Video Summarization Using Descriptors of Motion Activity: A Motion Activity based approach to Key-Frame Extraction from Video Shots" Journal of Electronic Imaging Oct. 2001 vol. 10 n 4 pp. 909-916.  
 Goncalves, "Towards a Learning Model for Feature Integration in Attention Control", retrieved on Jul. 19, 2006 at <<<http://ieeexplore.ieee.org/search/srchabstract.jsp?arnumber=1013553&isnumber=21824&pun=>>>>, International Conference on Multisensor Fusion and Integration for Intelligent Systems, Aug. 20-22, 2001, pp. 311-316.  
 Gu et al., "Dissolve detection in MPEG compressed video" IEEE pp. 1692-1696 1997.  
 Hargrove et al., "Logo Detection in Digital Video" Math 100 Mar. 6, 2001 <http://toonarchive.com/logo-detection/> 9 pages.  
 Held, "Focus on Agere System's Orinoco PC Card" International Journal of Network Management Jan. 2002 pp. 187-193.  
 Heng et al., "Post shot boundary detection technique: Flashlight scene determination" University of Western Australia pp. 447-450 1999.  
<http://www.fatalfx.com/nomad/> "NOMAD (No More Advertising)" Feb. 17, 2003 4 pages.

- Itti et al., "A Comparison of Feature Combination Strategies for Saliency-Based Visual Attention Systems: Proc. of SPIE Human Vision and Electronic Imaging IV (HVEI'99) San Jose CA vol. 3644 pp. 473-482 Jan. 1999.
- Itti et al., "A Model of Saliency-based Visual Attention for Rapid Scene Analysis" IEEE Trans. on Pattern Analysis and Machine Intelligence 1998 5 pages.
- Itti et al., "A Model of Saliency-based Visual Attention for Rapid Scene Analysis" IEEE Transactions on Pattern Analysis and Machine Intelligence 1998.
- Itti et al., "Computational Modelling of Visual Attention" Nature Reviews/Neuroscience vol. 2 Mar. 2001 pp. 1-11.
- Jha, "Wireless Landscape-Need for Seamless Connectivity" Wireless Personal Communications 2002 pp. 275-283.
- Jing et al., "An Effective Region-Based Image Retrieval Framework" From ACM Multimedia 2002 28 pages.
- Lee et al., "Automatic Video Parsing Using Shot Boundary Detection and Camera Operation Analysis" Pattern Recognition, Elsevier, Kidlington, GB, vol. 34, No. 3, Mar. 2001, pp. 711-719.
- Lee et al., "Perception-Based Image Transcoding for Universal Multimedia Access" School of Electrical Engineering Korea University Seoul Korea 2001 IEEE pp. 475-478.
- Lelescu et al., "Real-time scene change detection on compressed multimedia bitstream based on statistical sequential analysis" IEEE pp. 1141-1144 2000.
- Li, et al., "Statistical Learning of Multi-View Face Detection" Proceedings of ECCV 2002 25 pages.
- Lin, et al., "Video Scene Extraction by Force Competition" IEEE Intl. Conference on Multimedia and Expo (ICME 001) Waseda University Tokyo Japan Aug. 22-25, 2001 4 pages.
- Lu, et al., "A Robust Audio Classification and Segmentation Method" Microsoft Research China 9 pages.
- Lu, et al. "Content Analysis for Audio Classification and Segmentation" IEEE Transactions on Speech and Audion Processing vol. 10 No. 7 Oct. 2002 pp. 504-516.
- Lu, et al., "Content-Based Audio Segmentation Using Support Vector Machines" 2001 IEEE pp. 956-959.
- Ma, et al., "A Model of Motion Attention for Video Skimming" Microsoft Research Asia 4 pages.
- Ma, et al., "A New Perceived Motion Based Shot Content Representation" Microsoft Research China 4 pages.
- Ma, et al., "A User Attention Model for Video Summarization" Proceedings of ICIP 2002.
- Milanese et al., "Attentive Mechanisms for Dynamic and Static Scene Analysis" Optical Engineering v34 No. 8 Aug. 1995 pp. 2428-2434.
- Niebur et al., "Computational Architectures for Attention" The Attentive Brain Chapter 9 1998 pp. 163-186.
- "An MPEG-1 shot boundary detector using XIL colour histograms" Dublin City University pp. 1-7 1998.
- O'Toole et al., "Evaluation of Automatic Shot Boundary Detection on a Large Video Test Suite" School of Computer Applications & School of Electronic Engineering Dublin City University Glasnevin Dublin Ireland Challenge of Image Retrieval Newcastle 1999 pp. 1-12.
- Osberger, et al., "An Automatic Image Quality Assessment Technique Incorporating Higher Level Perceptual Factors", retrieved on Jul. 19, 2006, at <<<http://ieeexplore.ieee.org/search/srchabstract.jsp?arnumber=727227&isnumber=15680&punu>>>, International Conference on Image Processing, vol. 3, Oct. 4-7, 1998, pp. 414-418.
- Sadler, "Automatic TV Advertisement Detection from MPEG Bitstream" Centre for Digital Video Processing/Research Institute for Network & Communication Eng. Dublin City University 12 pages.
- Bennett, et al., "Proscenium: A Framework for Spatio-Temporal Video Editing", ACM, 2003, pp. 8.
- Fan, et al., "Visual Attention Based Image Browsing on Mobile Devices", IEEE, 2003, pp. 4.
- Foote, et al., "Creating Music Videos using Automatic Media Analysis", ACM, 2002, pp. 8.
- Girgensohn, et al., "A Semi-automatic Approach to Home Video Editing", ACM, 2000, pp. 9.
- Girgensohn, et al., "Home Video Editing Made Easy—Balancing Automation and User Control", pp. 8.
- Graham, et al., "The Video Paper Multimedia Playback System", ACM, 2003, pp. 2.
- Hua, et al., "Automatically Converting Photographic Series into Video", ACM, 2004, pp. 8.
- Hua, et al., "AVE—Automated Home Video Editing", ACM, 2003, pp. 8.
- Hua, et al., "Video Booklet", IEEE, 2005, pp. 4.
- Yip, et al., "The Automatic Video Editor", ACM, 2003, pp. 2.
- Adams, et al., "Seeded Region Growing", IEEE, vol. 16, No. 6, 1994, pp. 641-648.
- Alexander, et al., "Informedia: News-on-Demand Multimedia Information Acquisition and Retrieval", 1997, pp. 1-21.
- DeMenthon, et al., "Video Summarization by Curve Simplification", ACM, 1998, pp. 211-218.
- Gong, et al., "Creating Motion Video Summaries with Partial Audio-Visual Alignment", IEEE, 2002, pp. 285-288.
- Beg, "User Feedback Based Enhancement in Web Search Quality", Elsevier Inc., 2004, pp. 20.
- Degn, et al., "Peer Group Filtering and Perceptual Color Image Quantization", IEEE, 1999, pp. 8.
- Kelly, et al., "Implicit Feedback for Inferring User Preference: A Bibliography", SIGIR Forum, vol. 37, No. 2, 2003, pp. 18-28.
- "NOMAD (No More Advertising)", <http://www.fatalfx.com/nomad/>, Feb. 17, 2003 4 pages.
- Pirolli, "Exploring Browser Design Trade-offs Using a Dynamical Model of Optimal Information Foraging", Proc. SIGCHI Conference on Human Factors in Computing Systems CHI, 1998, pp. 33-40.
- Smith, et al., "Scalable Multimedia Delivery for Pervasive Computing", ACM, 1999, pp. 131-140.
- Sugiyama, et al., "Adaptive Web Search Based on User's Implicit Preference", DEWS, 2004, pp. 8.
- Rauschenbach, et al., "Demand-driven image transmission with levels of detail and regions of interest", Computers & Graphics, 1999, pp. 857-866.
- Rauschenbach, et al., "Flexible Embedded Image Communication using Levels of Detail and Regions of Interest", Proc. IMC'98, 1998, pp. 1-10.
- Gamaz et al., "Robust scene-change detection in MPEG compressed domain" Elec & Comp. Eng. pp. 95-99 1998.
- Lienhart, et al., "On the Detection and Recognition of Television Commercials" University of Mannheim 17 pages.
- Dufaux, "Key frame selection to represent a video", Proc. 2000 Int. Con. on Image Processing, vol. 2, 9-2000, pp. 275-278.

\* cited by examiner

100

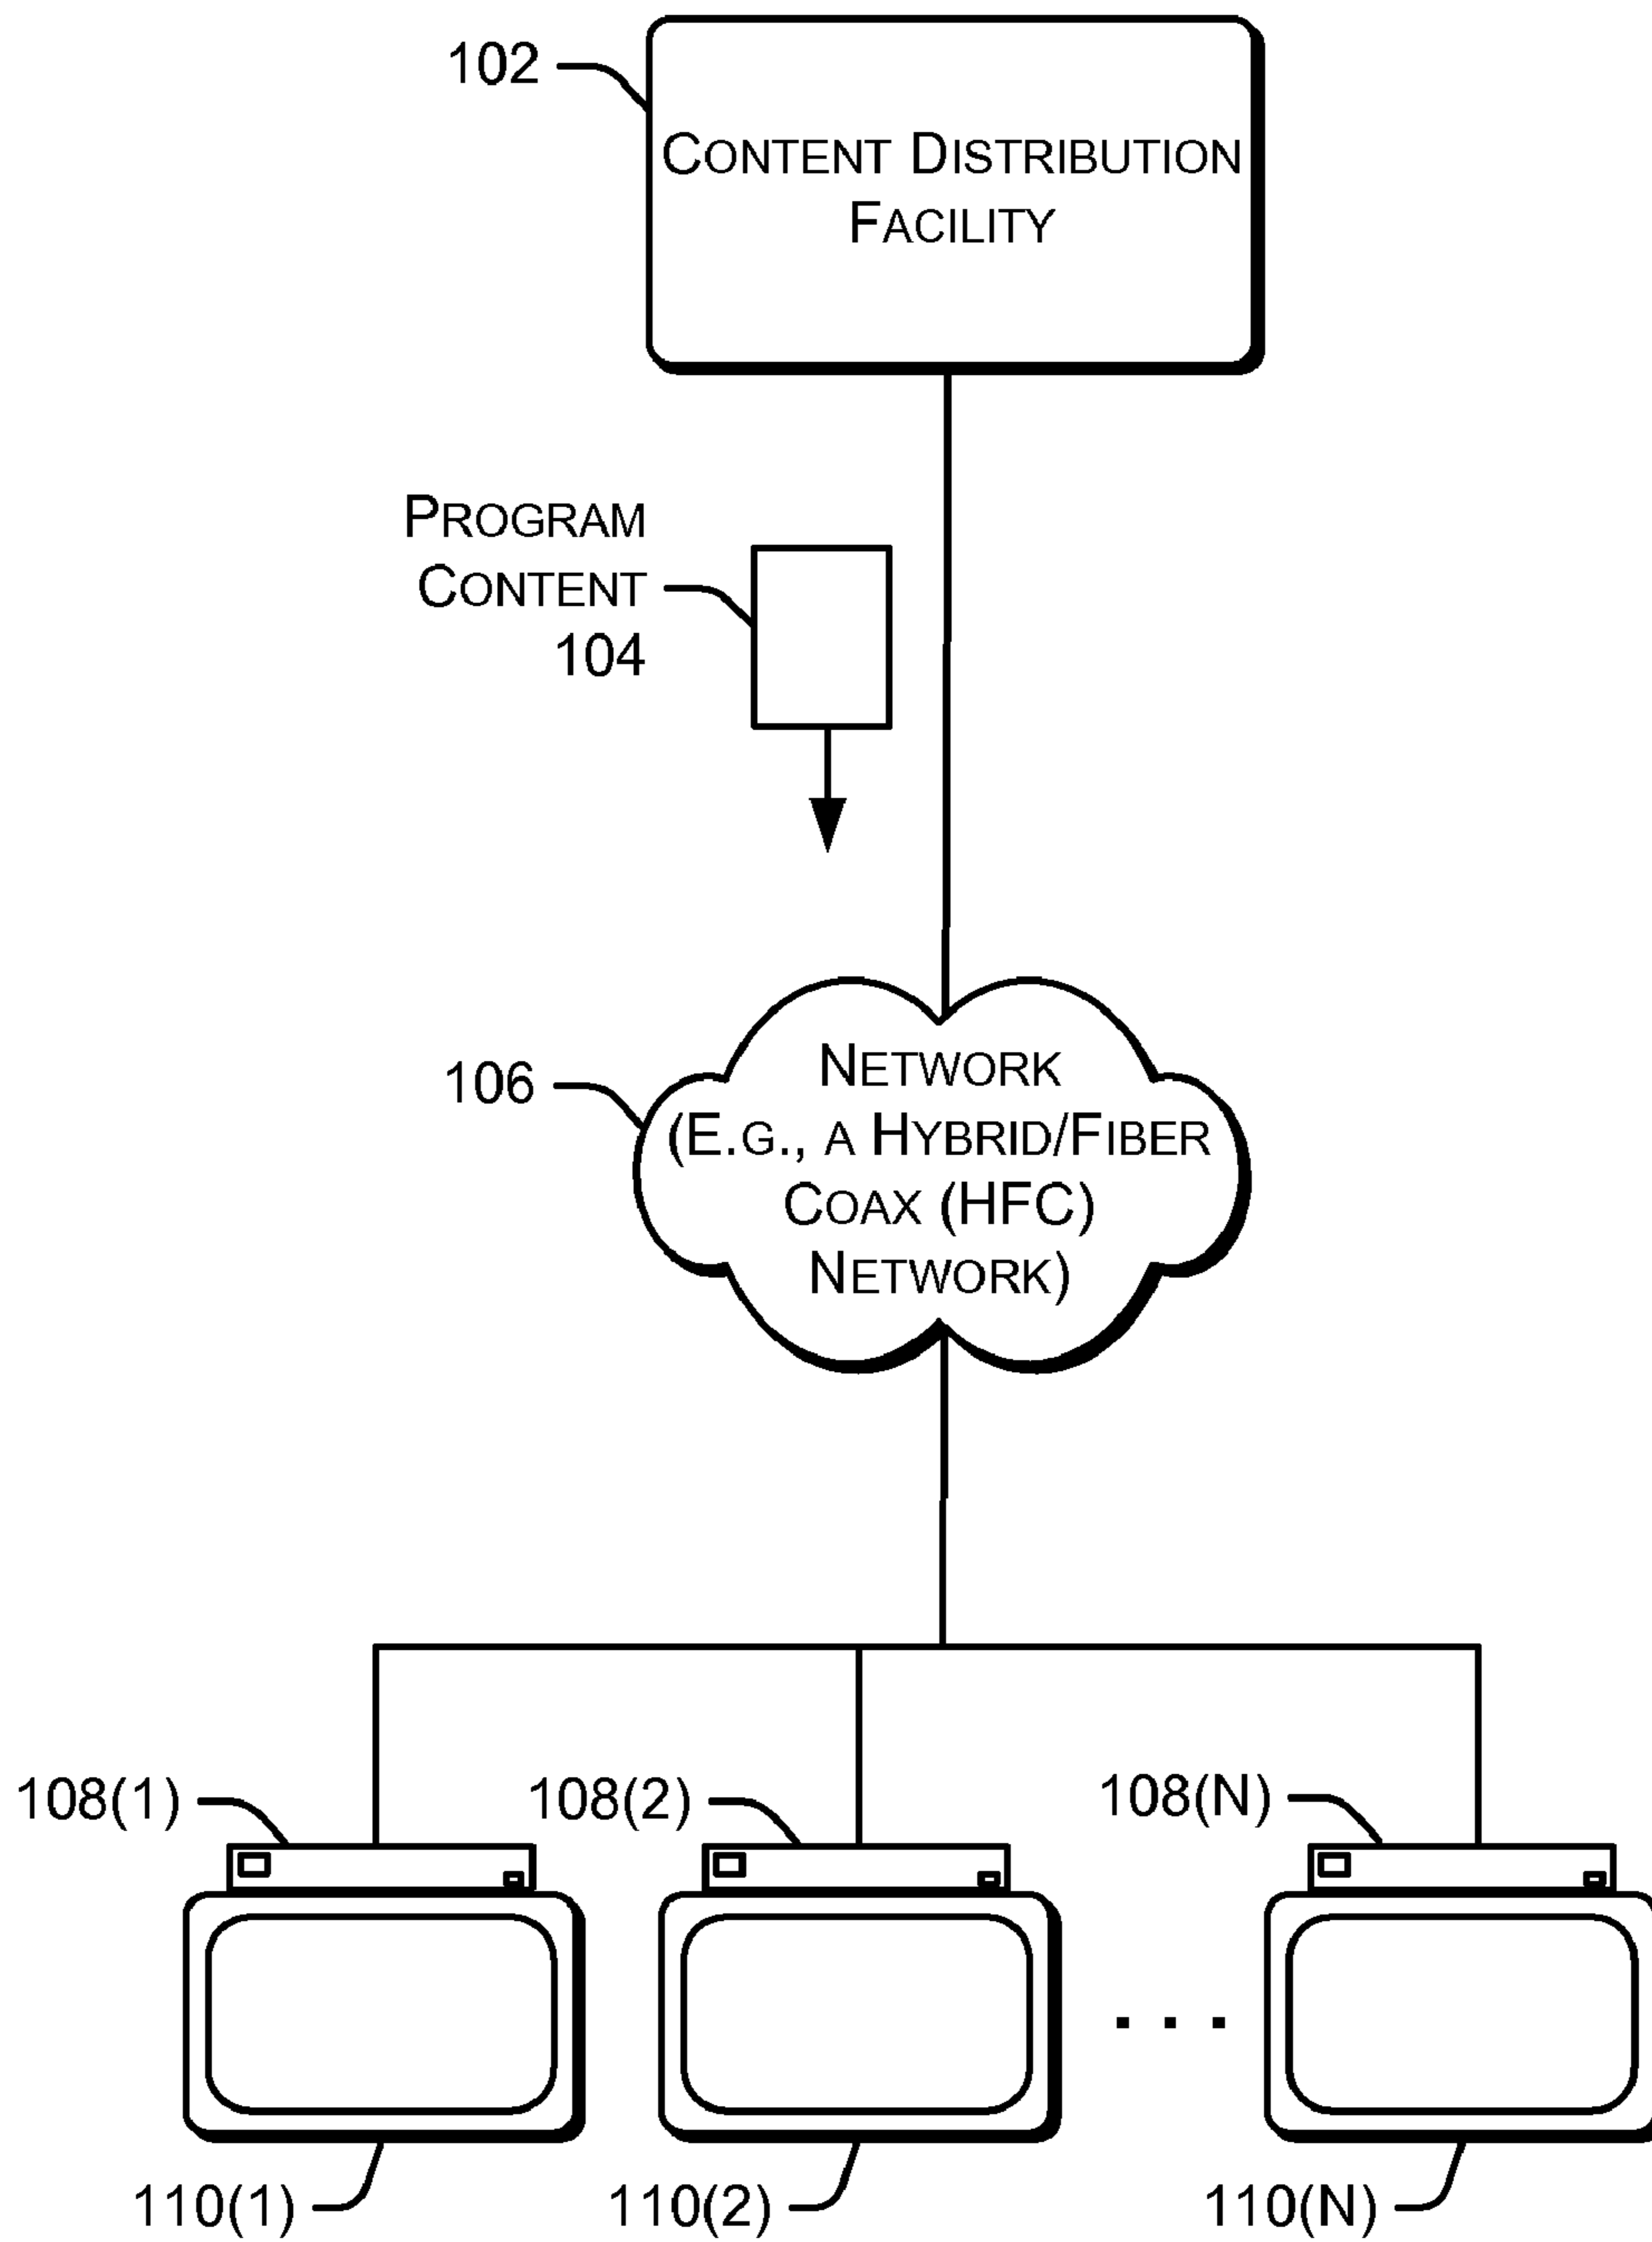


Fig. 1

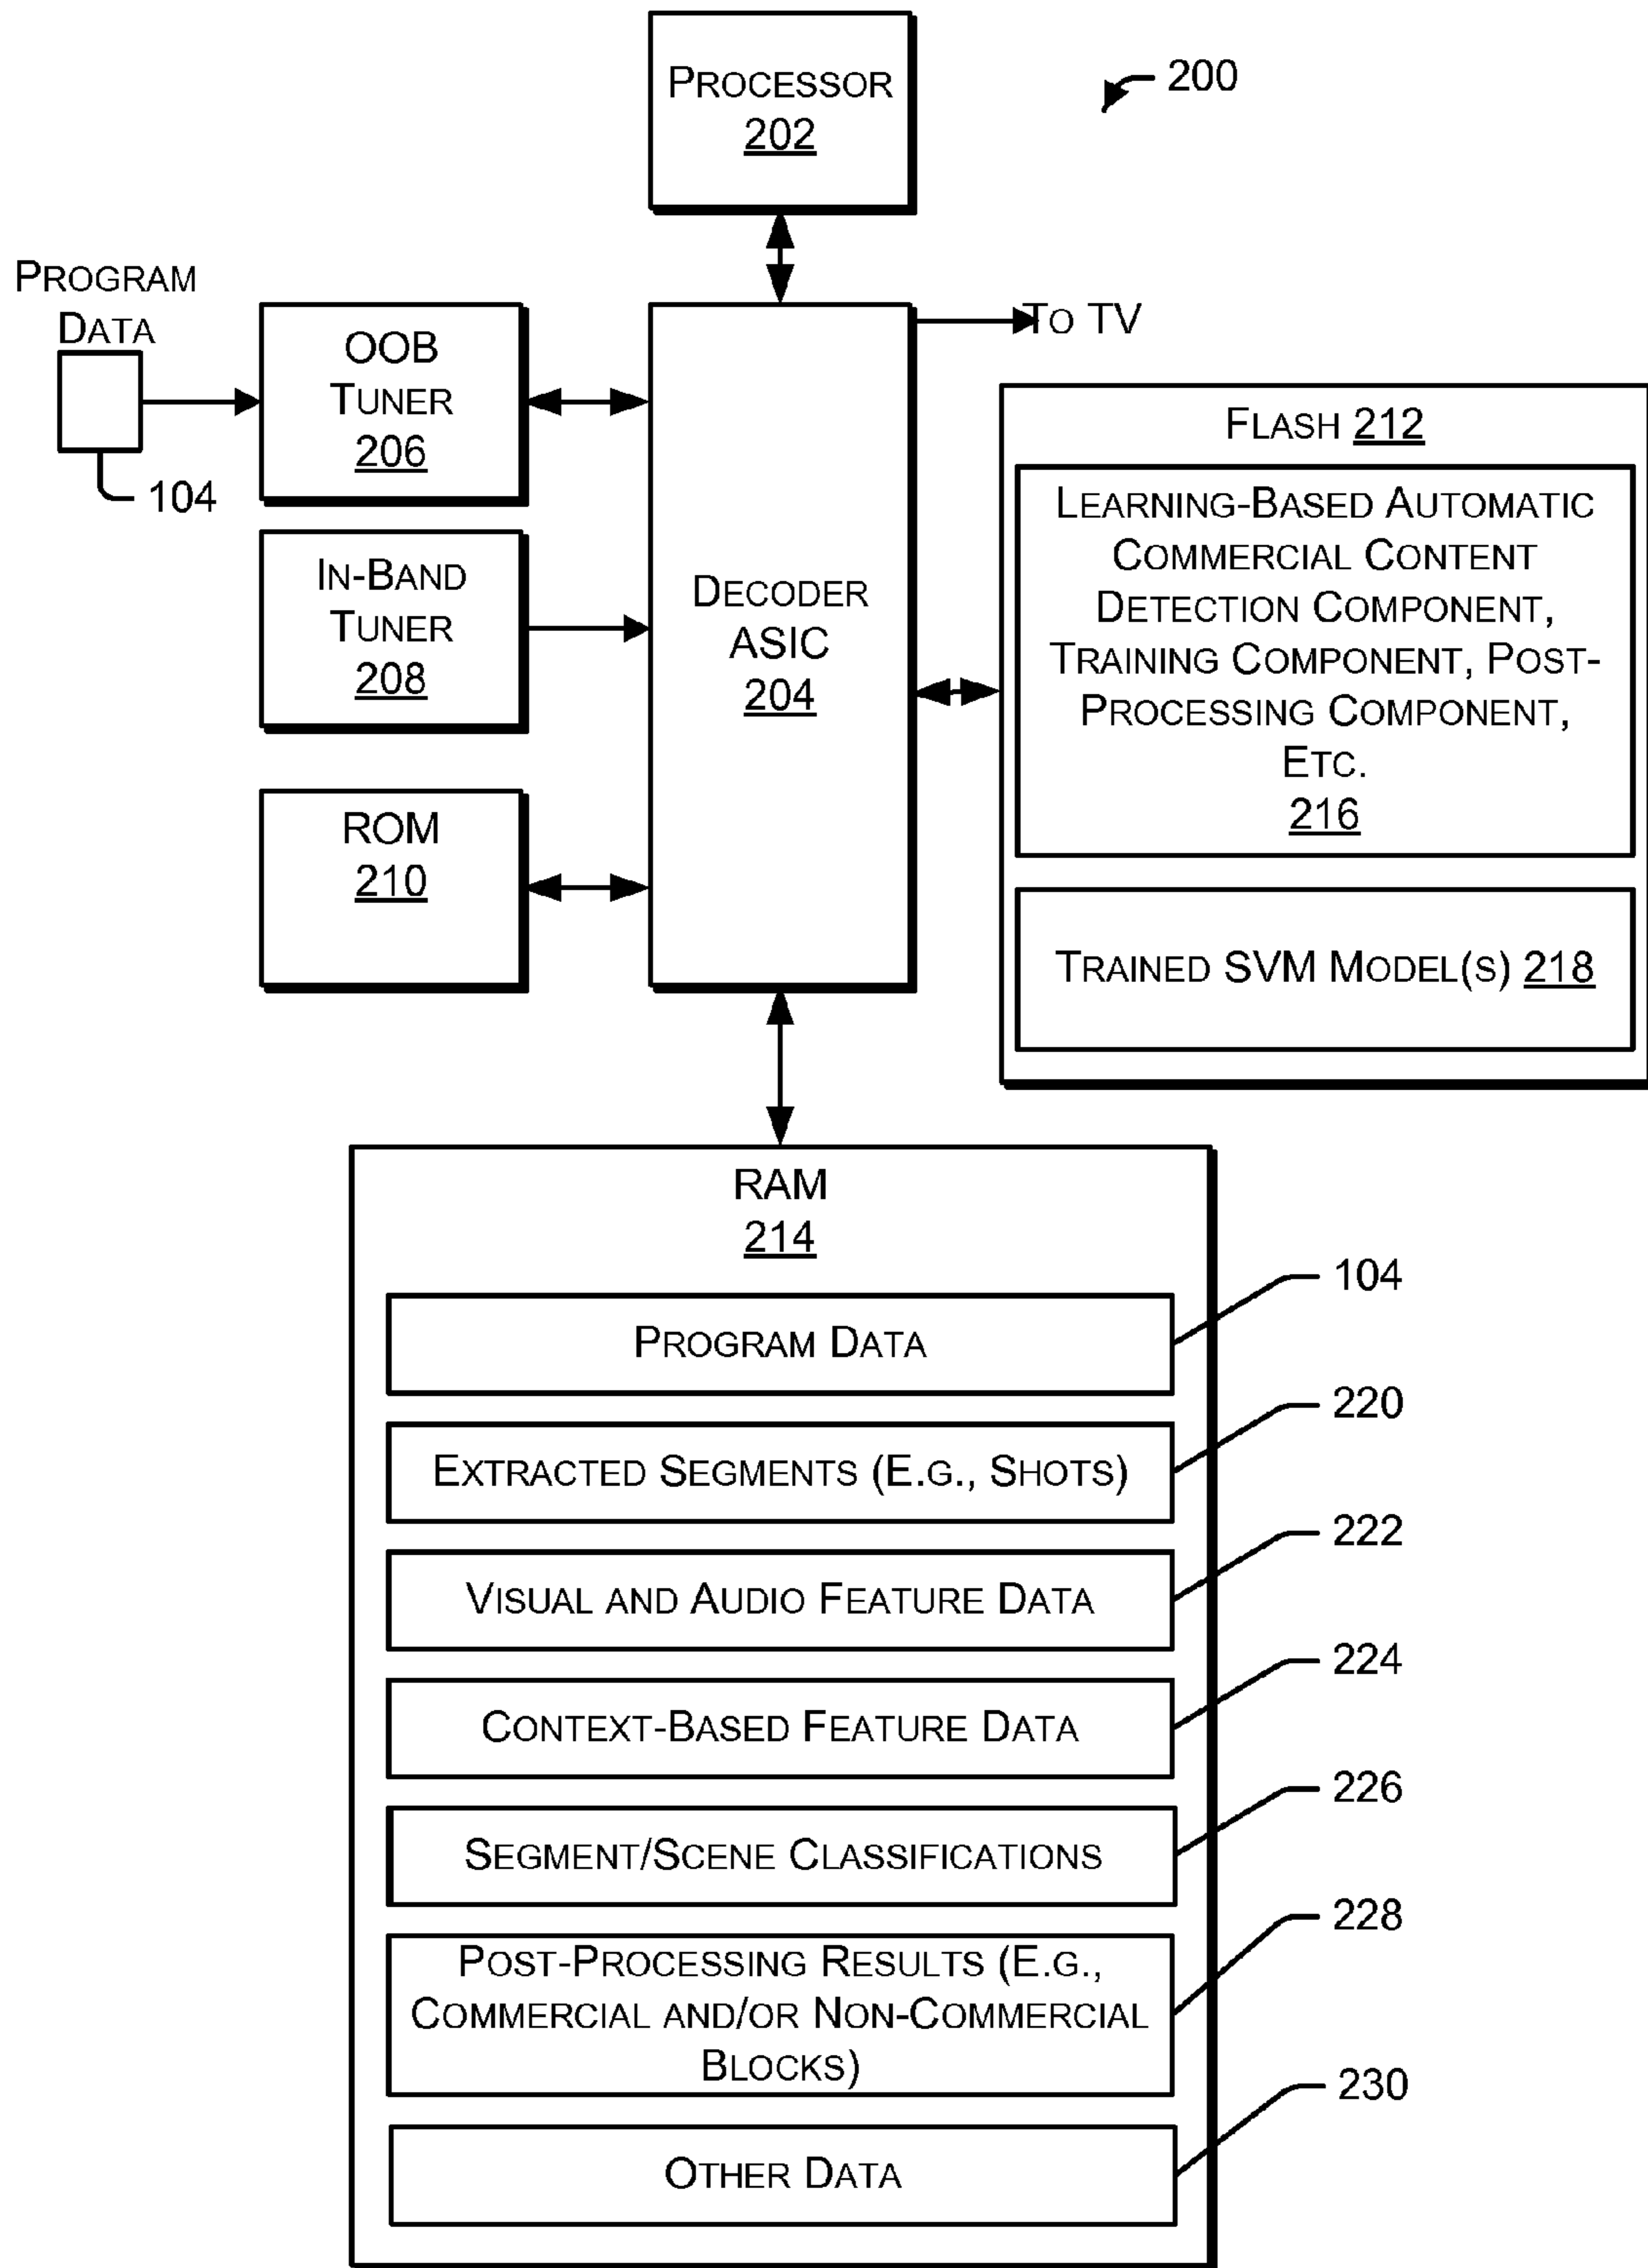
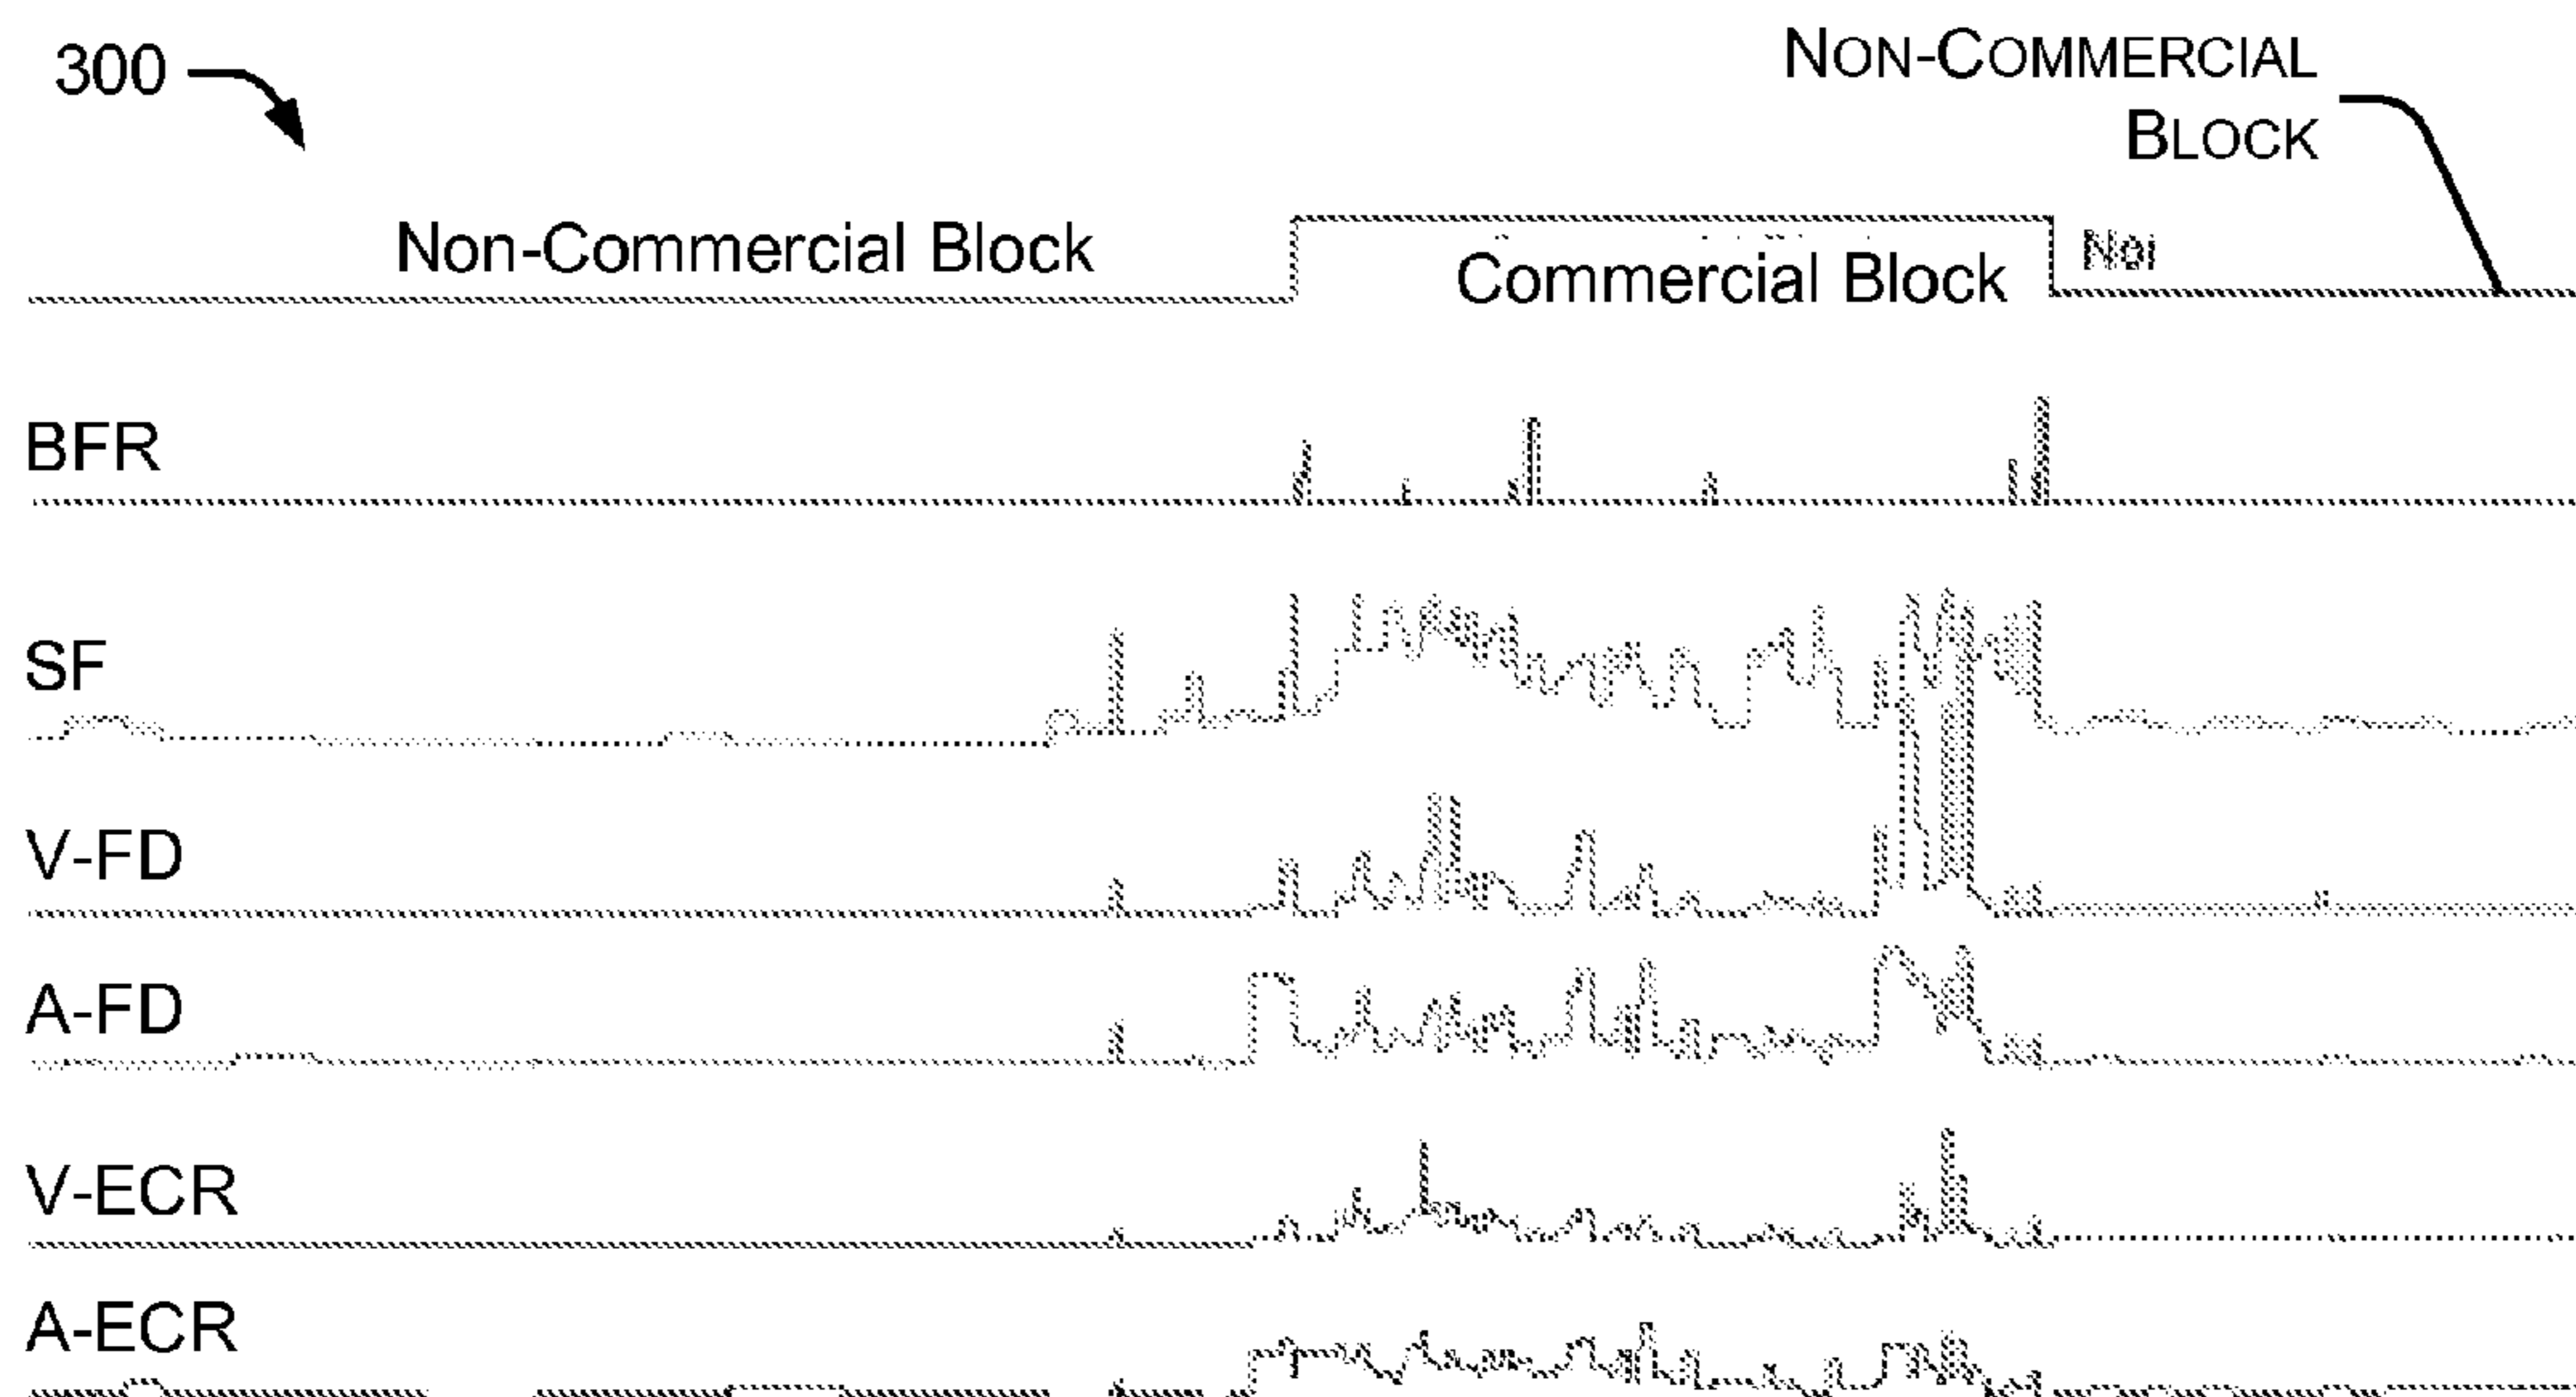
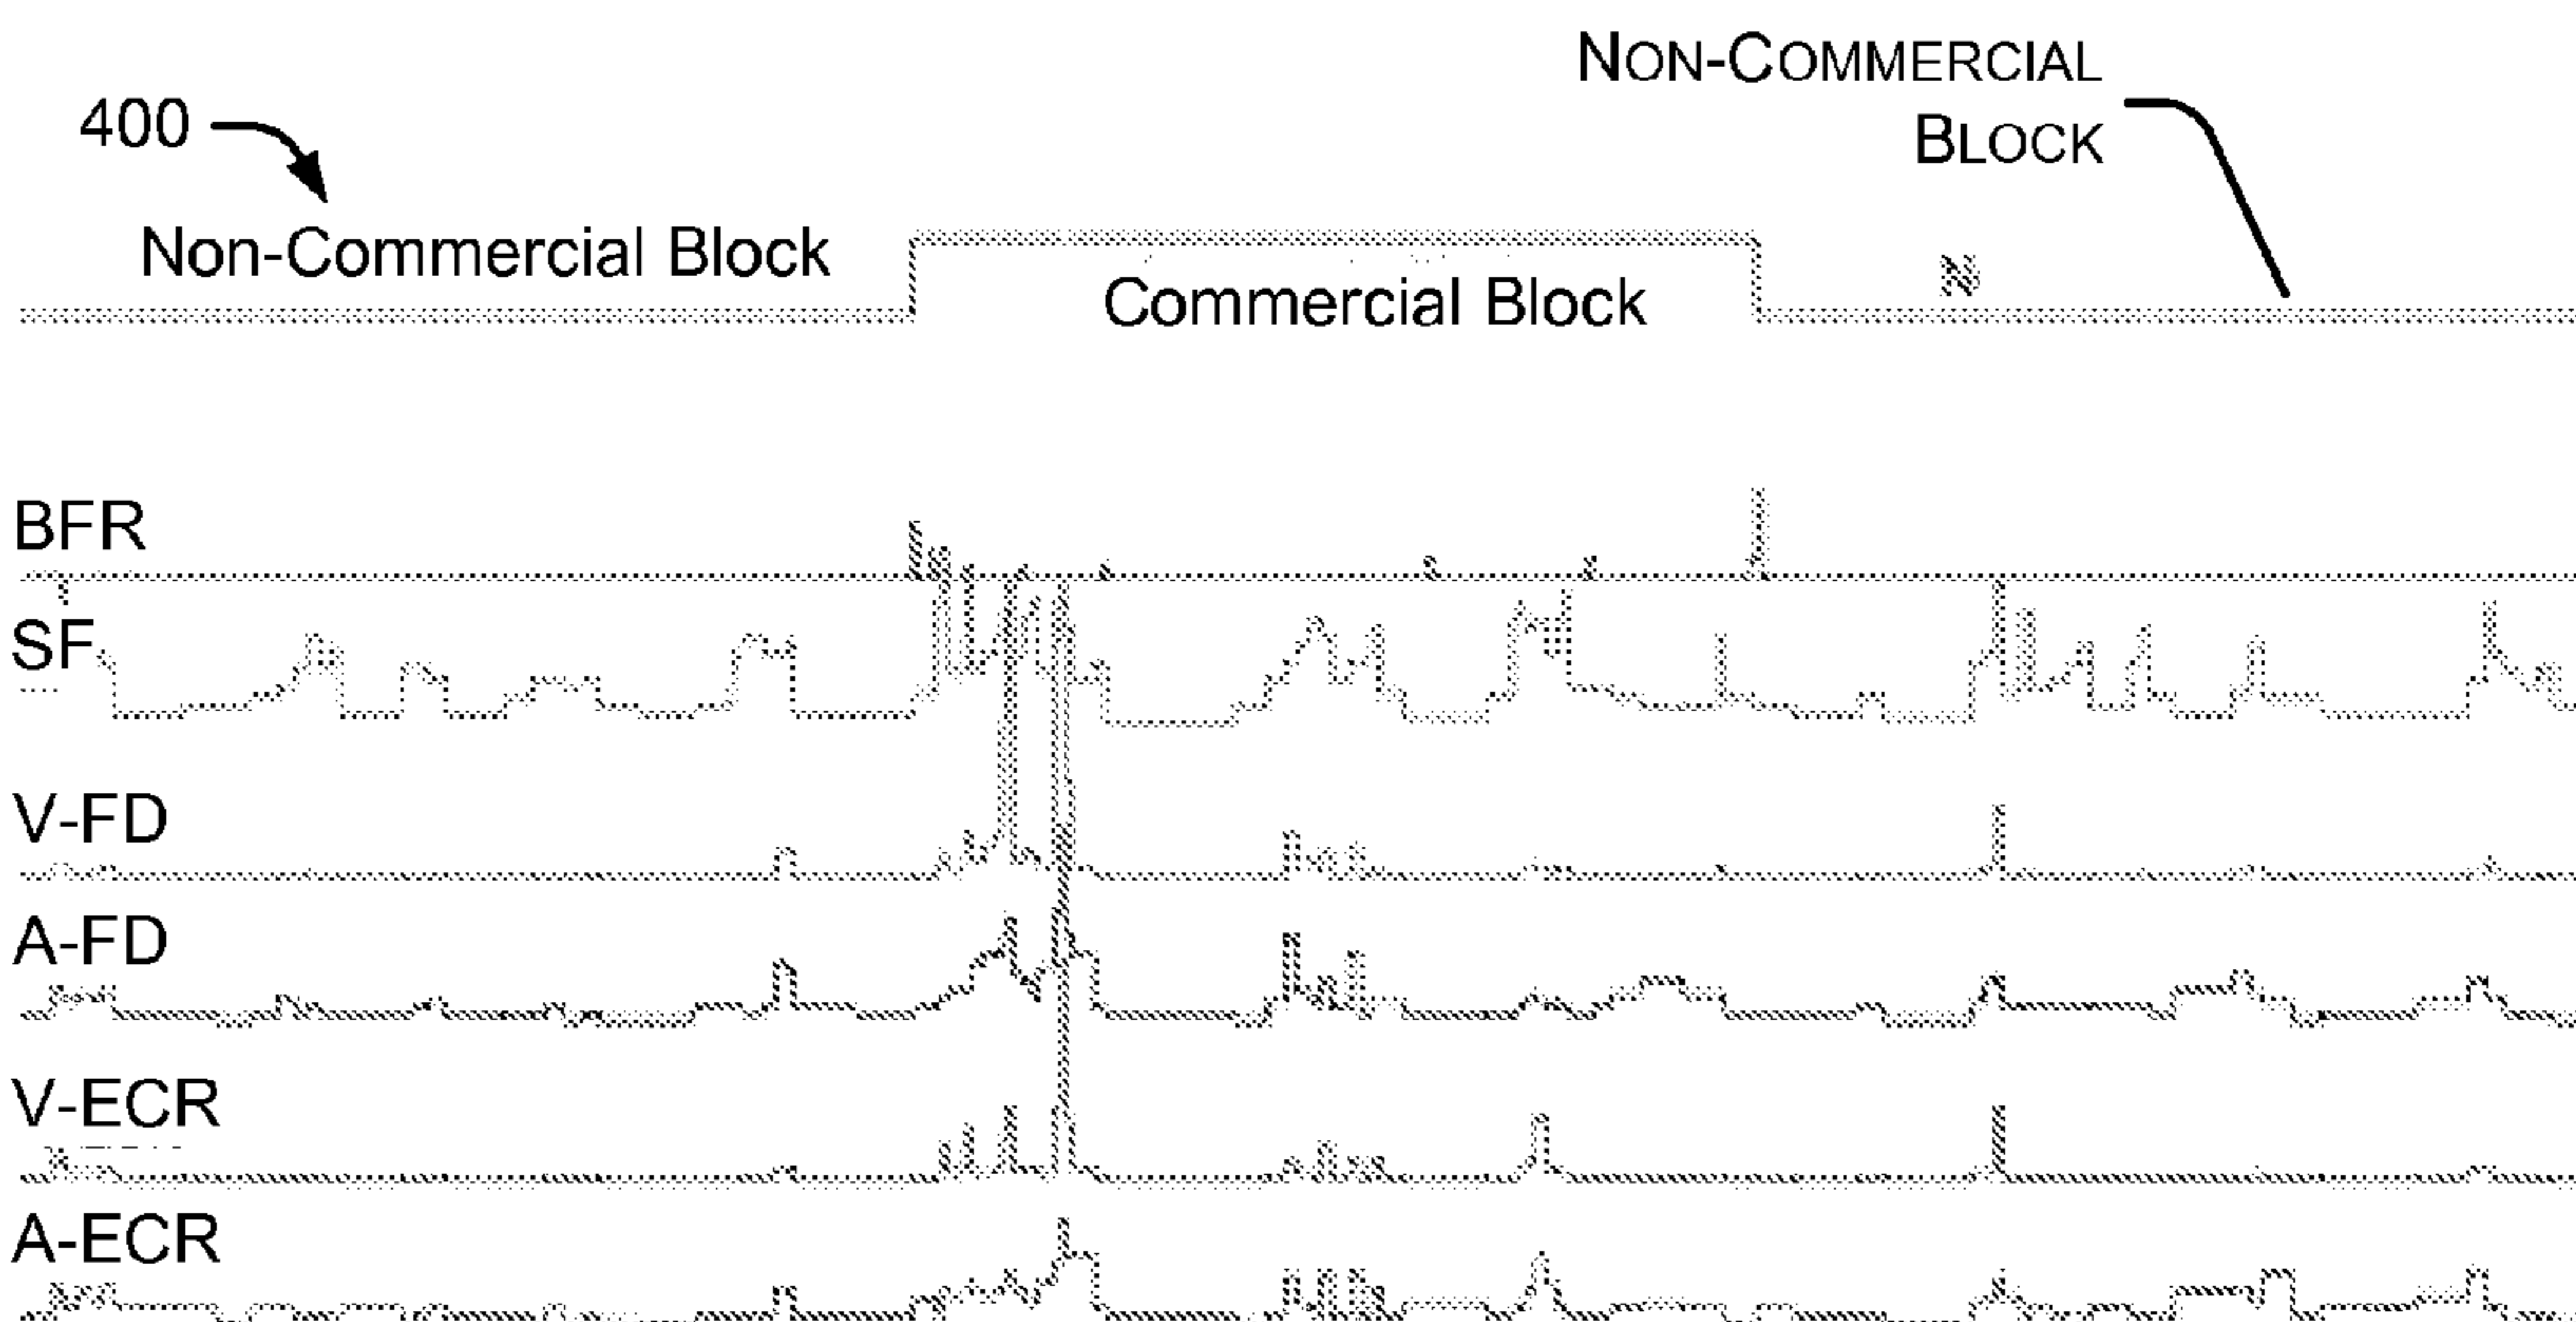


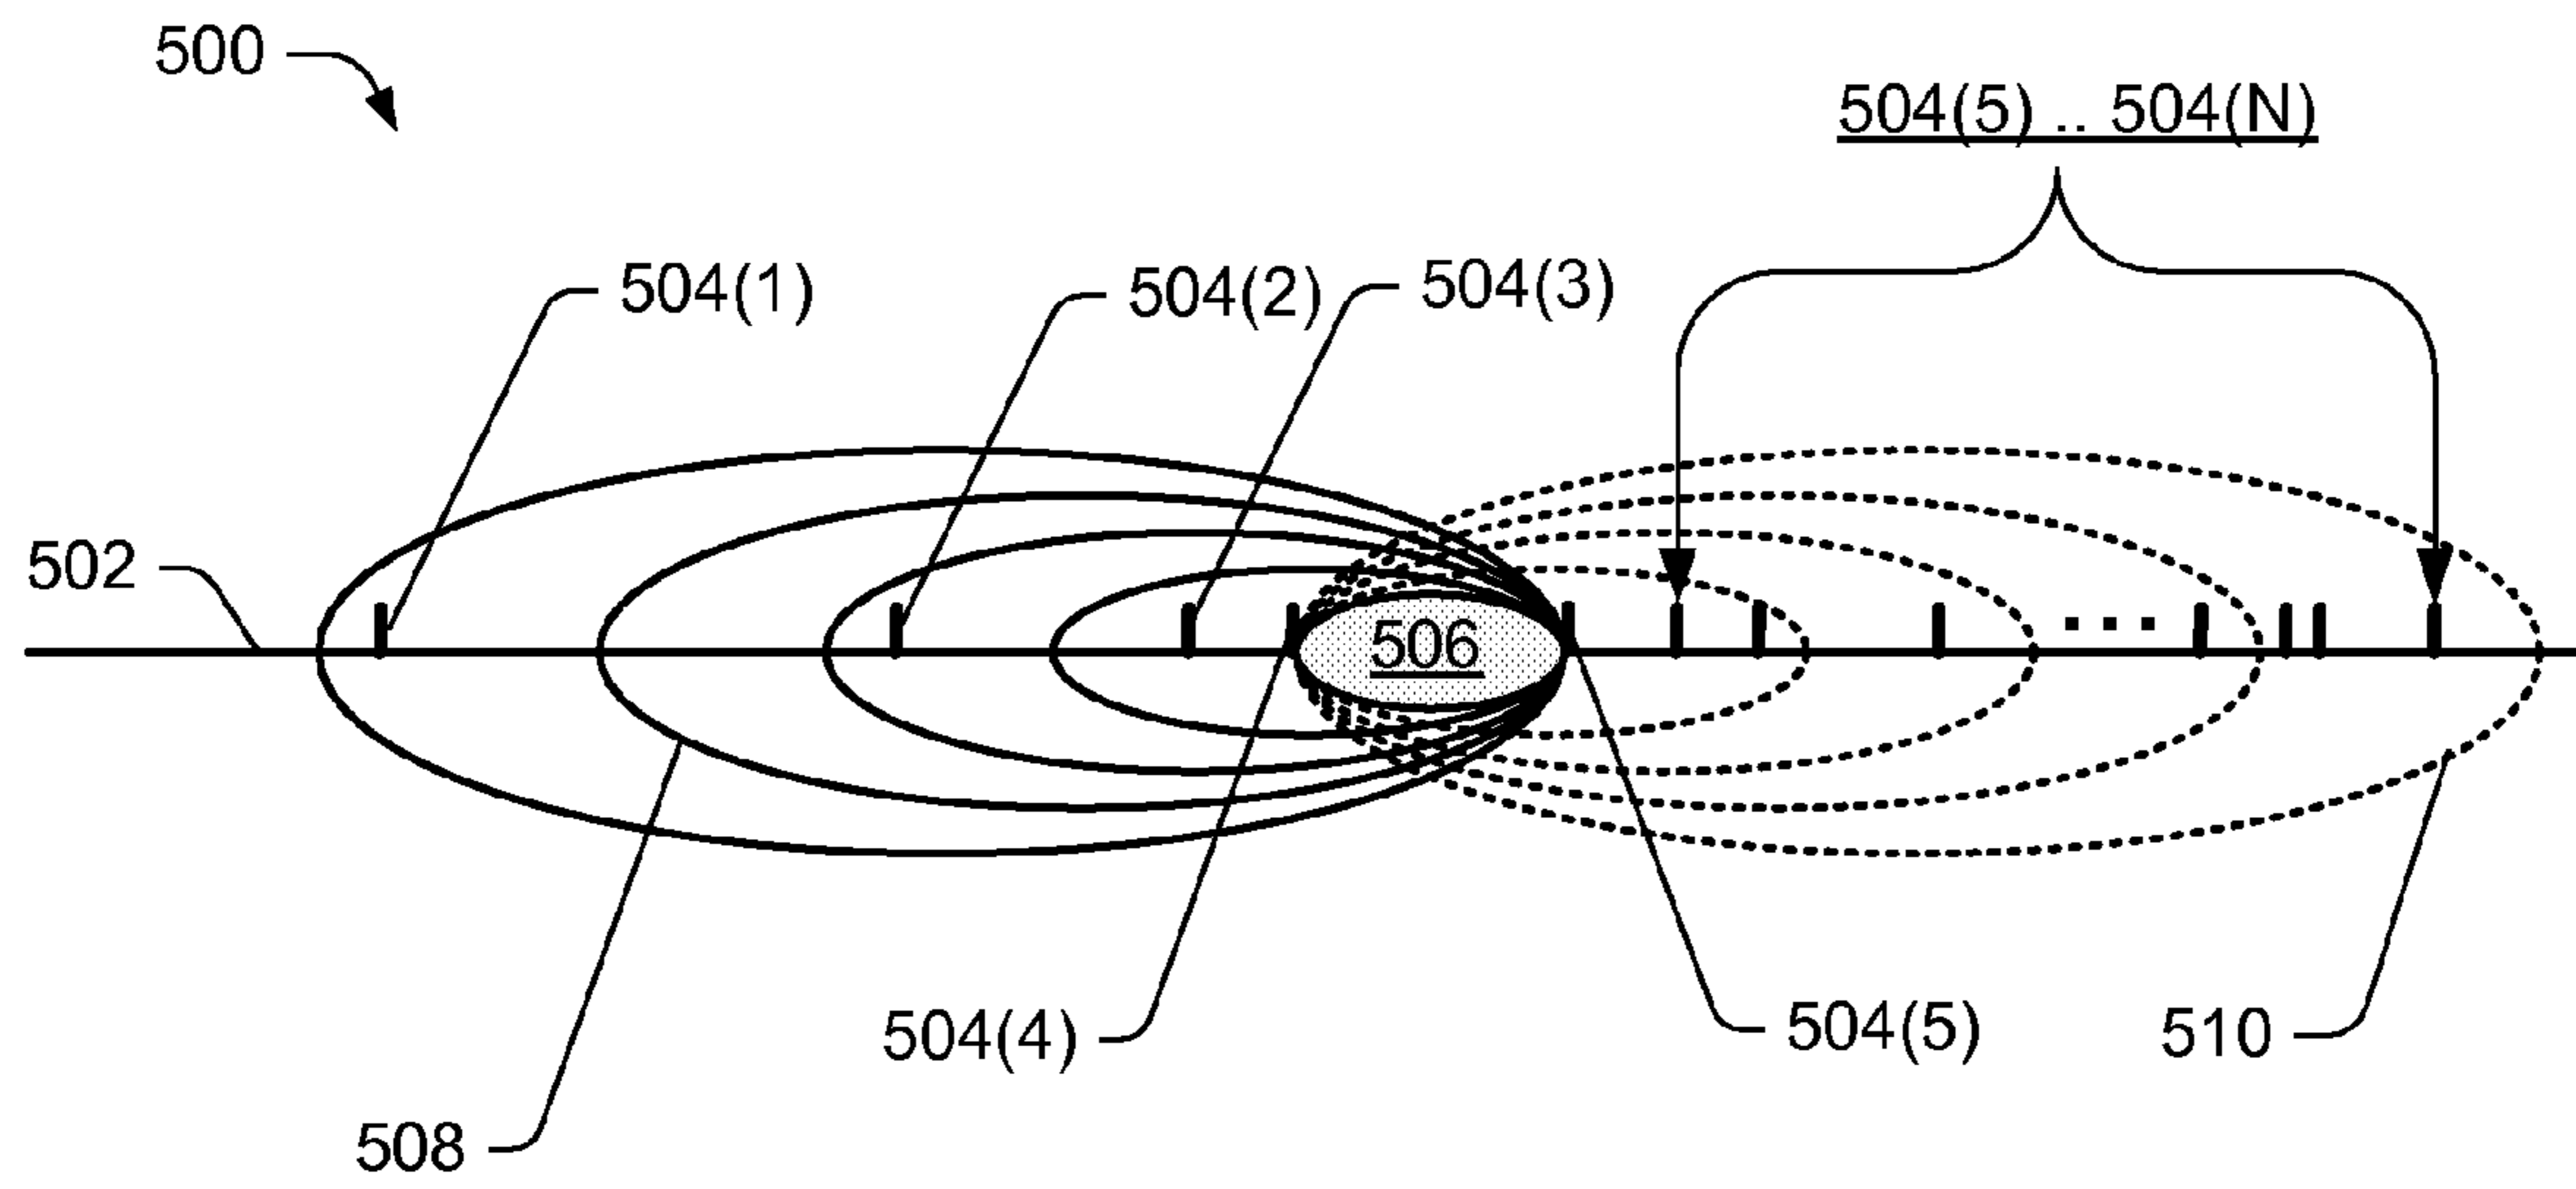
Fig. 2



*Fig. 3*



*Fig. 4*



*Fig. 5*

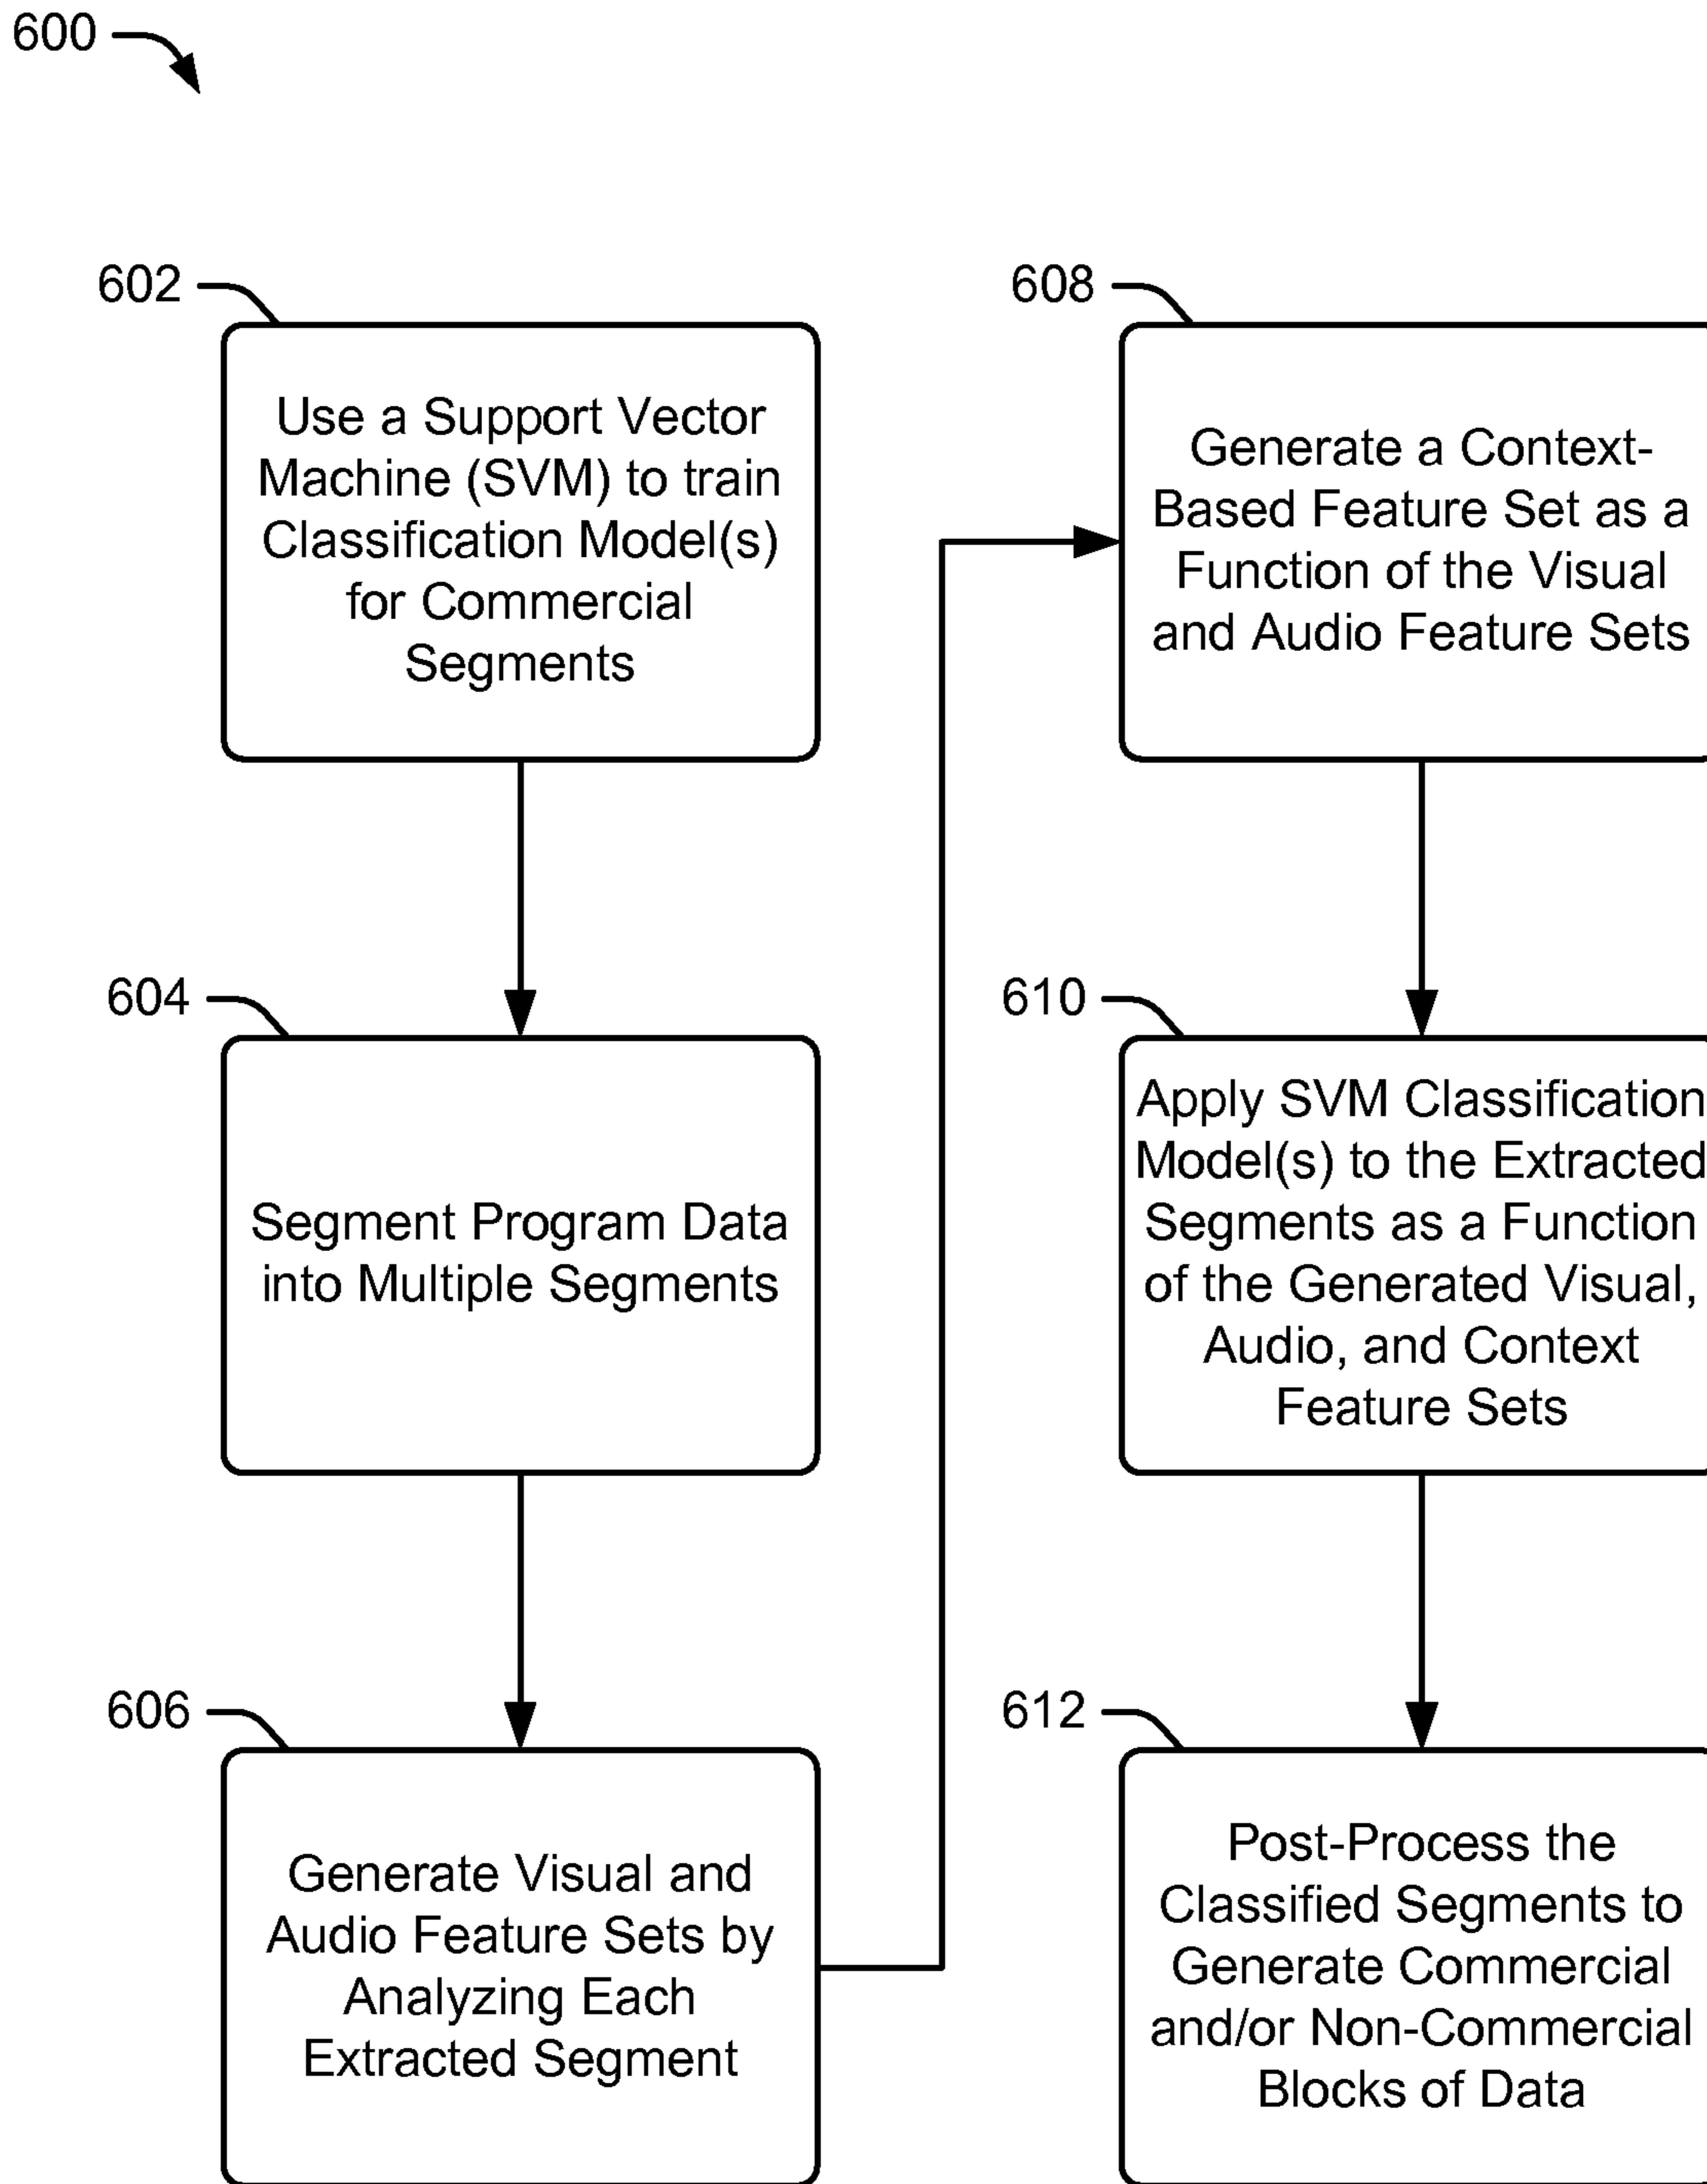
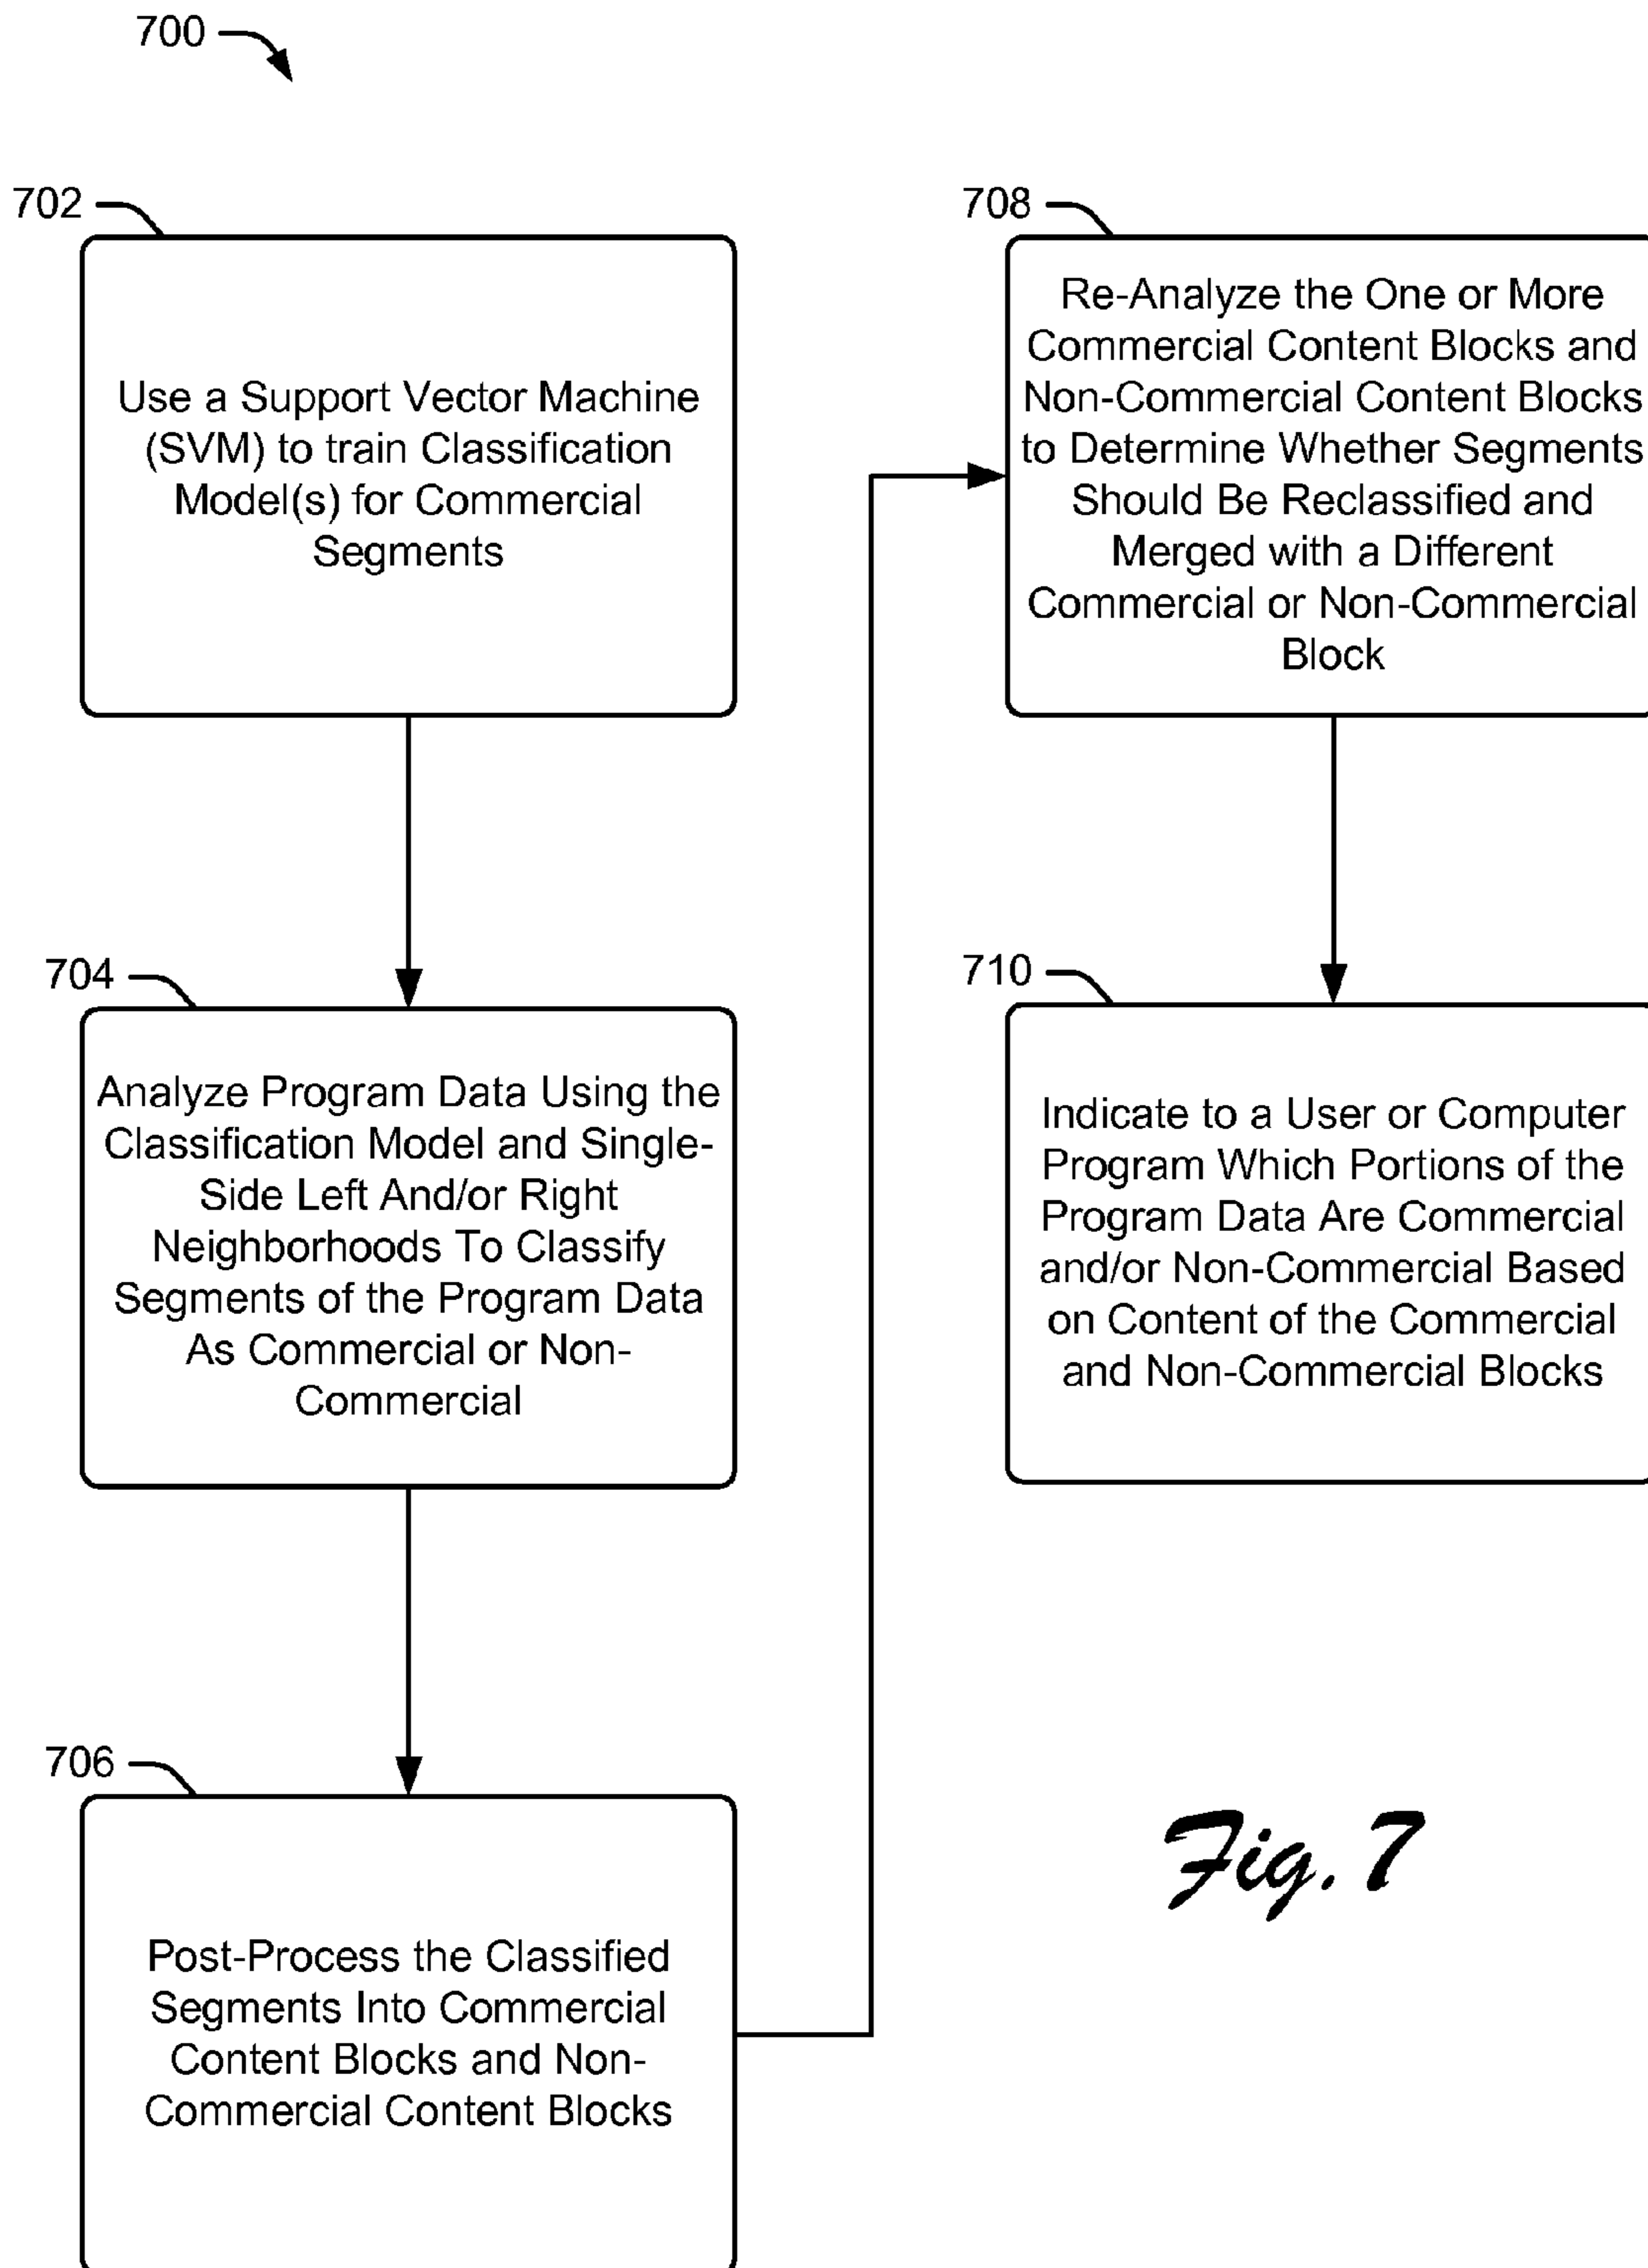


Fig. 6





*Fig. 7*

1

## LEARNING-BASED AUTOMATIC COMMERCIAL CONTENT DETECTION

### RELATED APPLICATION

This patent application is a continuation of U.S. patent application Ser. No. 10/368,235, titled "Learning-Based Automatic Commercial Content Detection", filed on Feb. 18, 2003, and hereby incorporated by reference.

### BACKGROUND

There are many objectives for detecting TV commercials. For example, companies who produce commercial advertisements (ads) generally charge other companies to verify that certain TV commercials are actually broadcast as contracted (e.g., broadcast at a specified level of quality for a specific amount of time, during a specific time slot, and so on). Companies who design ads typically research commercials to develop more influential advertisements. Thus, commercial detection techniques may also be desired to observe competitive advertising techniques or content.

Such commercial content verification/observation procedures are typically manually performed by a human being at scheduled broadcast time(s), or by searching (forwarding, rewinding, etc.) a record of a previous broadcast. As can be appreciated, waiting for a commercial to air (broadcast), setting up recording equipment to record a broadcast, and/or searching records of broadcast content to verify commercial content airing(s) can each be time consuming, laborious, and costly undertakings.

To make matters even worse, and in contrast to those that desire to view TV commercials, others may find commercial content aired during a program to be obtrusive, interfering with their preferred viewing preferences. That is, rather than desiring to view commercial content, such entities would rather not be presented with any commercial content at all. For example, a consumer may desire to record a TV program without recording commercials that are played during broadcast of the TV program. Unfortunately, unless a viewer actually watches a TV program in its entirety to manually turn on and off the recording device to selectively record non-commercial content, the viewer will typically not be able to record only non-commercial portions of the TV program.

In light of the above, whether the objective is to view/record commercial content or to avoid viewing/recording commercial content, existing techniques for commercial content detection to enable these goals are substantially limited in that they can be substantially time consuming, labor intensive, and/or largely ineffective across a considerable variety of broadcast genres. Techniques to overcome such limitations are greatly desired.

### SUMMARY

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the detailed description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. In view of this, systems and methods for learning-based automatic commercial content detection are described. In one aspect, the systems and methods include a training component and an analyzing component. The training component trains a commercial content classification model using a kernel support vector machine. The analyzing component analyzes pro-

2

gram data such as video and audio data using the commercial content classification model and one or more of single-side left neighborhood(s) and right neighborhood(s) of program data segments. Based on this analysis, each of the program data segments are classified by the systems and methods as being commercial or non-commercial segments. Further aspects of the systems and methods for learning-based automatic commercial content detection are presented in the following sections.

### BRIEF DESCRIPTION OF THE DRAWINGS

The following detailed description is described with reference to the accompanying figures. In the figures, the left-most digit of a component reference number identifies the particular figure in which the component first appears.

FIG. 1 shows an exemplary computing environment on which systems, apparatuses and methods for learning-based automatic commercial content detection may be implemented, according to an embodiment.

FIG. 2 is a block diagram of an exemplary client computing device that includes computer-readable media with computer-program instructions for execution by a processor to implement learning-based automatic commercial content detection, according to an embodiment.

FIG. 3 shows results of an exemplary application of time and segment-based visual criteria to extracted segments of a digital data stream (program data such as a television program) to differentiate non-commercial content from commercial content, according to an embodiment.

FIG. 4 shows results of an exemplary application of visual criteria to extracted segments of a digital data stream, wherein visual feature analysis by itself does not conclusively demarcate the commercial content from the non-commercial content, according to an embodiment.

FIG. 5 is an exemplary diagram showing single-sided left and right "neighborhoods" of a current segment (e.g., shot) that is being evaluated to extract a context-based feature set, according to an embodiment.

FIG. 6 shows an exemplary procedure for implementing learning-based automatic commercial content detection, according to an embodiment.

FIG. 7 shows an exemplary procedure for implementing learning-based automatic commercial content detection, according to an embodiment.

### DETAILED DESCRIPTION

#### Overview

The following discussion is directed to systems and methods for learning-based automatic detection of commercial content in program data. Program data, for example, is data that is broadcast to clients in a television (TV) network such as in interactive TV networks, cable networks that utilize electronic program guides, Web-enabled TV networks, and so on. Program data may also be embodied as digital video and/or audio data that has been stored onto any of numerous different types of volatile or non-volatile data storage such as computer-readable media, tapes, CD-ROMs, diskettes, and so on. Numerous computing architectures such as in a set-top box, a digital program recorder, or a general purpose PC can be modified according to the following description to practice learning-based automatic detection of commercial content.

To this end, a commercial content classification model is trained using a kernel support vector machine (SVM). The model is trained with commercial content that represents any

number of visual and/or audio genres. Techniques to generate SVM-based classification models are well known.

If the trained SVM model is not generated on the particular computing device that is to implement the learning-based automatic commercial content detection operations, the trained SVM model is manually or programmatically uploaded or downloaded to/from the particular computing device. For example, if the trained SVM model is posted to a Web site for download access, any number of client devices (e.g., set-top boxes, digital recorders, etc.) can download the trained model for subsequent installation. For purposes of this discussion, the particular computing device is a digital recorder that has either been manufactured to include the trained SVM model or downloaded and installed the trained SVM model, possibly as an update.

At this point and responsive to receiving program data, the digital recorder (DR) divides the program data into multiple segments using any of numerous segmenting or shot boundary determination techniques. In one implementation, the multiple segments may be the same as shots. However, the segments, or segment boundaries are independent of shot boundaries and thus, do not necessarily represent shots. Rather, the segments are window/blocks of data that may or may not represent the boundaries of one or more respective shots. The DR analyzes the extracted segments with respect to multiple visual, audio, and context-based features to generate visual, audio, and context-based feature sets. The DR then evaluates the extracted segments in view of the trained SVM model to classify each of the extracted segments as commercial or non-commercial content. These classifications are performed in view of the visual, audio, and context-based feature sets that were generated from the extracted segments.

The DR then performs a number of post-processing techniques to provide additional robustness and certainty to the determined segment (e.g., shot) classifications. Such post-processing techniques include, for example, scene-grouping and merging to generate commercial and non-commercial blocks of content. As part of post-processing operations, these generated blocks are again evaluated based on multiple different threshold criteria to determine whether segments within blocks should be reclassified, merged with a different block, and/or the like.

In this manner, the systems, apparatus, and methods for learning-based commercial content detection differentiate and organize commercial and non-commercial portions of program data. Since the differentiated portions have been aggregated into blocks of like-classified segments/scenes, an entity that desires to verify/observe only commercial portions may do so without experiencing time consuming, labor intensive, and/or potentially prohibitive expenses that are typically associated with existing techniques. Moreover, an entity that desires to view/record only non-commercial portions of program data may do so without watching a program in its entirety to manually turn on and off the recording device to selectively record only non-commercial content.

#### An Exemplary System

Turning to the drawings, wherein like reference numerals refer to like elements, the invention is illustrated as being implemented in an exemplary computing environment. The exemplary computing environment is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of systems and methods the described herein. Neither should the exemplary computing environment be interpreted as having

any dependency or requirement relating to any one or combination of components illustrated in the computing environment.

FIG. 1 shows exemplary computing environment **100** on which systems, apparatuses and methods for learning-based automatic commercial content detection may be implemented. The exemplary environment represents a television broadcast system that includes a content distributor **102** for broadcasting program data **104** across network **106** to one or more clients **108(1)-108(N)**. The program data is broadcast via “wireless cable”, digital satellite communication, and/or other means. As used herein, program data refers to the type of broadcast data that includes commercial advertisements. The network includes any number and combination of terrestrial, satellite, and/or digital hybrid/fiber coax networks.

Clients **108(1)** through **108(N)** range from full-resource clients with substantial memory and processing resources (e.g., TV-enabled personal computers, multi-processor systems, TV recorders equipped with hard-disks) to low-resource clients with limited memory and/or processing resources (e.g., traditional set-top boxes, digital video recorders, and so on). Although not required, client operations for learning-based automatic commercial detection are described in the general context of computer-executable instructions, such as program modules stored in the memory and being executed by the one or more processors. Program modules generally include routines, programs, objects, components, data structures, etc., that perform particular tasks or implement particular abstract data types.

In one implementation, a client **108** (any one or more of clients **108(1)** through **108(N)**) is coupled to or incorporated into a respective television viewing device **110(1)** through **110(N)**.

#### An Exemplary Client

FIG. 2 shows an exemplary client computing device (i.e., one of clients **108(1)** through **108(N)** of FIG. 1) that includes computer-readable media with computer-program instructions for execution by a processor to implement learning-based automatic commercial content detection. For purposes of discussion, the exemplary client is illustrated as a general-purpose computing device in the form of a set-top box **200**. The client **200** includes a processor **202** coupled to a decoder ASIC (application specific integrated circuit) **204**. In addition to decoder circuitry, ASIC **204** may also contain logic circuitry, bussing circuitry, and a video controller. The client **200** further includes an out-of-band (OOB) tuner **206** to tune to the broadcast channel over which the program data **104** is downloaded. One or more in-band tuners **208** are also provided to tune to various television signals. These signals are passed through the ASIC **204** for audio and video decoding and then to an output to a television set (e.g., one of TVs **110(1)** through **110(N)** of FIG. 1). With the tuners and ASIC **204**, the client is equipped with hardware and/or software to receive and decode a broadcast video signal, such as an NTSC, PAL, SECAM or other TV system video signal and provide video data to the television set.

One or more memories are coupled to ASIC **204** to store software and data used to operate the client **200**. In the illustrated implementation, the client has read-only memory (ROM) **210**, flash memory **212**, and random-access memory (RAM) **214**. One or more programs may be stored in the ROM **210** or in the flash memory **212**. For instance, ROM **210** stores an operating system (not shown) to provide a run-time environment for the client. Flash memory **212** stores a learning-based (LB) commercial detection program module **216** that is executed to detect commercial portions of the program

data **104**. Hereinafter, the LB automatic commercial content detection program module is often referred to as “LBCCD” **216**. The LBCCD utilizes one or more trained SVM models **218**, which are also stored in the flash memory **214**, to assist in classifying portions of the program data **104** as commercial verses non-commercial.

RAM **214** stores data used and/or generated by the client **200** during execution of the LBCCD module **216**. Such data includes, for example, program data **104**, extracted segments **220**, visual and audio feature data **222**, context-based feature **224**, segment/scene classifications **226**, post processing results **228**, and other data **230** (e.g., a compression table used to decompress the program data). Each of these program module and data components are now described in view of the exemplary operations of the LBCCD module **216**.

To detect commercial portions of program data **104**, the LBCCD module **216** first divides program data **104** into multiple segments (e.g., shots). These segments are represented as extracted segments **220**. Segment extraction operations are accomplished using any of a number of known segmentation/shot extraction techniques such as those described in “A New Shot Detection Algorithm” D. Zhang, W. Qi, H. J. Zhang, 2nd IEEE Pacific-Rim Conference on Multimedia (PCM2001), pp. 63-70, Beijing, China, October 2001.

In another implementation, the LBCCD module **216** detects shot boundaries for program data shot extraction using techniques described in U.S. patent application Ser. No. 09/882,787, titled “A Method and Apparatus for Shot Detection”, filed on Jun. 14, 2001, commonly assigned herewith, and which is hereby incorporated by reference.

#### Time-based and Segment-based Visual Feature Analysis

The LBCCD module **216** analyzes each extracted segment **220** with respect to numerous visual and audio features to generate visual and audio feature data **222**. Although the extracted segments can be evaluated in view of any number of visual and audio criteria, in this implementation, six (6) visual features and five (5) audio features are used to analyze the extracted segments. Two (2) of the visual features are time-based features, and four (4) of the visual features are segment-based features. The time-based visual features include, for example, shot frequency (SF) and black frame rate (BFR) of every second.

Each segment is evaluated with respect to the segment-based visual features, which include Average of Edge Change Ratio (“A-ECR”), Variance of Edge Change Ratio (“V-ECR”), Average of Frame Difference (“A-FD”), and Variance of Frame Difference (“V-FD”). Edge Change Ratio represents the amplitude of edge changes between two frames [6] as follows:

$$ECR_m = \max\left(\frac{X_m^{in}}{\sigma_m}, \frac{X_{m-1}^{out}}{\sigma_{m-1}}\right). \quad (1)$$

Variable  $\sigma_m$  is the number of edge pixels in frame  $m$ ,  $X_m^{in}$  and  $X_{m-1}^{out}$  are the number of entering and exiting edge pixels in frame  $m$  and  $m-1$ , respectively. A-ECR and V-ECR of segment  $C$  are defined as

$$AECR(C) = \frac{1}{F-1} \sum_{m=1}^{F-1} ECR_m, \quad (2)$$

and

-continued

$$VECR(C) = \frac{1}{F-1} \sum_{m=1}^{F-1} (ECR_m - AECR(C))^2, \quad (3)$$

where  $F$  is the number of frames in the segment.

Frame Difference (FD) is defined by

$$FD_m = \frac{1}{P} \sum_{i=0}^{P-1} |F_i^m - F_i^{m-1}| \quad (4)$$

where  $P$  is the pixel number in one video frame,  $F_i^m$  is the intensity value of pixel  $i$  of frame  $m$ , and A-FD and V-FD are obtained similarly to A-ECR and V-ECR.

FIG. 3 shows results of an exemplary application of time and segment-based visual criteria to extracted segments of a digital data stream (e.g., program data such as a television program) to differentiate non-commercial content from commercial content. In this example, the horizontal axis of the graph **300** represents the passage of time, and the vertical axis of the graph **300** represents the values of respective ones of the calculated visual features as a function of time. The illustrated visual feature values include calculated A-ECR, V-ECR, A-FD, V-FD, BFR, and SF results.

As shown, the visual feature values of graph **300** clearly distinguish a commercial block of program data from a non-commercial block. However, how clearly such content can be differentiated across different portions of the program content using only such visual feature calculations is generally a function of the visual attributes of the program data at any point in time. Thus, depending on visual feature content of the program data, visual feature analysis by itself may not always be able to clearly demarcate commercial portions from non-commercial portions. An example of this is shown in FIG. 4, wherein visual feature analysis by itself does not conclusively demarcate the commercial content from the non-commercial content.

In light of this, and to add additional content differentiation robustness to LBCCD **216** (FIG. 2) algorithms, the LBCCD module further generated audio and context-based feature sets for each of the extracted segments **220** (FIG. 2). These additional feature sets provide more data points to verify against the trained SVM model(s) **218** (FIG. 2) as described below.

#### 50 Audio Features

Referring to FIG. 2, the LBCCD module **216** further analyzes each of the extracted segments **220** (FIG. 2) with respect to audio break frequency and audio type. In this implementation, audio break frequency and audio type determinations are extracted from the segments at constant time intervals, for example, every  $\frac{1}{2}$  second. With respect to audio break detection, audio transitions are typically present between commercial and non-commercial or different commercial program data. Such audio breaks are detected as a function of speaker change, for example, as described in L. Lu, H. J. Zhang, H. Jiang, “Content Analysis for Audio Classification and Segmentation” IEEE Trans on Speech and Audio Processing, Vol. 10, No. 7, pp. 504-516, October 2002, and which is incorporated by reference.

For instance, the LBCCD module **216** first divides the audio stream from the program data **104** into sub-segments delineated by a sliding-window. In one implementation, each

sliding-window is three (3) seconds wide and overlaps any adjacent window(s) by two-and-one-half (2½) seconds. The LBCCD module further divides the sub-segments into non-overlapping frames. In one implementation, each non-overlapping frame is twenty-five (25) ms long. Other window and sub-segment sizes can be used and may be selected according to numerous criteria such as the genre of the program data, and so on. At this point, The LBCCD module **216** extracts Mel-frequency Cepstral Coefficient (MFCC) and short-time energy from each non-overlapping frame. K-L distance is used to measure the dissimilarity of MFCC and energy between every two sub-segments,

$$D(i, j) = \frac{1}{2}tr[(C_i - C_j)(C_j^{-1} - C_i^{-1})] + \frac{1}{2}tr[(C_i^{-1} + C_j^{-1})(u_i - u_j)(u_i - u_j)^T].$$

This is equation (5), wherein  $C_i$  and  $C_j$  are the estimated covariancematrixes,  $u_i$  and  $u_j$  are the estimated mean vectors, from  $i$ -th and  $j$ -th sub-segment respectively; and  $D(i, j)$  denote the distance between the  $i$ -th and  $j$ -th audio sub-segments.

Thus, an audio transition break is found between  $i$ -th and  $(i+1)$ -th sub-segments, if the following conditions are satisfied:

$$D(i, i+1) > D(i+1, i+2), D(i, i+1) > D(i-1, i), D(i, i+1) > Th_i \quad (5)$$

The first two conditions guarantee that a local dissimilarity peak exists, and the last condition can prevent very low peaks from being detected.  $Th_i$  is a threshold, which is automatically set according to the previous  $N$  successive distances. That is:

$$Th_i = \alpha \cdot \frac{1}{N} \sum_{n=0}^N D(i-n-1, i-n) \quad (6)$$

wherein  $\alpha$  is a coefficient amplifier.

With respect to audio type discrimination, existing techniques typically utilize only silence (e.g., see [2] in APPENDIX) to determine whether there may be a break between commercial and non-commercial content. Such existing techniques are substantially limited in other indicators other than lack of sound (silence) can be used to detect commercial content. For example, commercial content typically includes more background sound than other types of programming content. In light of this, and in contrast to existing techniques, the LBCCD module **216** utilizes audio criteria other than just silence to differentiate commercial content from non-commercial content.

In this implementation, the LBCCD module **216** utilizes four (4) audio types, speech, music, silence and background sound, to differentiate commercial portion(s) of program data **104** from non-commercial portion(s) of program data. The LBCCD module analyzes each of the extracted segments **220** as a function of these audio types using techniques, for example, as described in "A Robust Audio Classification and Segmentation Method," L. Lu, H. Jiang, H. J. Zhang, 9th ACM Multimedia, pp. 203-211, 2001, and/or "Content-based Audio Segmentation Using Support Vector Machines," L. Lu, Stan Li, H. J. Zhang, Proceedings of ICME 2001, pp. 956-959, Tokyo, Japan, 2001, both of which are hereby incorporated by reference (see, [8] and [9] in APPENDIX).

Based on such audio type analysis, the LBCCD module **216** calculates a respective confidence value for each audio

type for each segment of the extracted segments **220**. Each confidence value is equal to the ratio of the duration, if any, of a specific audio type in a particular sub-segment.

#### Context-based Features

Without forehand knowledge of the content of a particular program, it is typically very difficult to view a program for only one or two seconds and based on that viewing, identify whether a commercial block or a non-commercial block was viewed. However, after watching some additional number of seconds or minutes, the viewed portion can typically be recognized as being commercial, non-commercial, or some combination of both (such as would be seen during a commercial/non-commercial transition).

The LBCCD module **216** takes advantage of the time-space relationship that enables one to identify context, wherein time is the amount of time that it takes to comprehend context, and wherein space is the amount of the program data **104** that is evaluated in that amount of time. In particular, the LBCCD module **216** identifies context-based information of a current segment from the current segment as well as from segments within single-sided neighborhoods of the current segment. Neighbor segments are other ones of the segments in extracted segments **220** that border the current segment at some distance to the left or right sides of the current segment. The size or distance of a neighborhood is a function of the ordinal number of the neighborhood.

FIG. **5** is an exemplary diagram **500** showing single-sided left and right "neighborhoods" of a current segment that is being evaluated to extract a context-based feature set. Horizontal axis **502** represents a sequence of program data **104** (FIGS. **1** and **2**). Vertical tick marks **504(1)** through **504(N)** represent program data segment boundaries. Each adjacent tick mark pair represents a boundary of a particular segment. Although some numbers of segment boundaries are shown, the actual number of segment boundaries in the program data can be just about any number since it will typically be a function of program data content and the particular technique(s) used to segment the program data into respective segments.

For purposes of discussion, segment **506** (the shaded oval) is selected as an exemplary current segment. The current segment is one of the extracted segments **220** (FIG. **2**). Segment boundaries **504(4)** and **504(5)** delineate the current segment. As the LBCCD module **216** (FIG. **2**) evaluates program data (represented by the horizontal axis **502**) to differentiate commercial from non-commercial content, each other segment in the program data is, at one time or another, designated to be a current segment to determine its context-based features.

Neighborhood(s) to the left of the right-most boundary of the exemplary current segment **506** are represented by solid line non-shaded ovals. For example, left neighborhood oval **508** encapsulates segment boundaries **504(2)**-**504(5)**. Neighborhood(s) to the right of the left-most boundary of the exemplary current segment are represented by dotted line non-shaded ovals. For example, right neighborhood oval **510** encapsulates segment boundaries **504(4)**-**504(N)**. As the respective left and right neighborhoods show, the LBCCD module **216** (FIG. **2**) evaluates segments in "single side" neighborhoods, each of which extend either to the left or to the right of the current segment's boundaries.

A single-side neighborhood is not a two-side neighborhood. This means that a neighborhood does not extend both to the left and to the right of a current segment's boundaries. This single-side aspect of the neighborhoods reduces undes-

ired “boundary effects” that may otherwise increase commercial/non-commercial classification errors at the boundaries of the extracted segments.

Referring to FIG. 2, the LBCCD module **216** generates context-based feature set **224** according to the following. Let  $[s_i, e_i]$  denote the start and end frame number of current segment  $C_i$ .  $[s_i, e_i]$  also represents start and end times (in seconds) of the segment (time-based features are a function of time). The  $(2n+1)$  neighborhoods include left  $n$  neighborhoods, right  $n$  neighborhoods, and the current segment  $C_i$ , are determined as follows:

$$N^k = [N_s^k, N_e^k] = \begin{cases} [\min(e_j + \alpha k, 0), e_i] & k < 0 \\ [s_i, e_i] & k = 0 \\ [s_i, \min(s_i + \alpha k, L)] & k > 0. \end{cases} \quad (7)$$

The variable  $L$  is the length or total frame number of the TV program, and  $k \in Z, |k| \leq n, Z$  is the set of integers,  $\alpha$  is the time step of the neighborhoods. In one implementation,  $n=6$ , and  $\beta=5$ .

Let  $S^k$  represent the set of all segments that are partially or totally included in  $N^k$ , that is

$$S^k = \{C_j^k : 0 \leq j < M^k\} = \{C_i : C_i \cap N^k \neq \Phi\} \quad (8)$$

where  $M^k$  is the number of segments in  $S^k$ ,  $i$  and  $j$  are non-negative integers.

Derived context-feature set **224** is the average value of basic features on  $S^k$  (for segment-based features) or  $N^k$  (for time-based features). For example, A-ECR on  $S^k$  and BFR on  $N^k$  are obtained by

$$AECR_{S^k} = \frac{1}{\sum_{j=0}^{M^k-1} (e_j^k - s_j^k)} \sum_{j=0}^{M^k-1} (e_j^k - s_j^k) AECR(C_j^k), \quad (9)$$

and

$$BFR_{N^k} = \frac{1}{N_e^k - N_s^k} \sum_{j=N_s^k}^{N_e^k-1} BFR(j), \quad (10)$$

where

$$[e_j^k, s_j^k]$$

is the start and end frame number of segment  $C_j^k$ . BFR( $j$ ) represents the black frame rate in  $[j, j+1]$  (count by second). Also,  $AECR_{S^k}$  is not equal to the average ECR in

$$[N_s^k, N_e^k]$$

Thus, ECR is not counted between two consecutive segments.

Using these techniques and in this implementation, the LBCCD model **216** generates  $11 \times (2n+1)$  context-based features **224** from the above eleven (11) described visual and audio features. Accordingly, in this implementation, the context-based feature set represents a one-hundred-and-forty-three (143) dimensional feature for each extracted segment **220**.

#### SVM-based Classification

To further classify individual ones of the extracted segments **220** as consisting of commercial or non-commercial content, the LBCCD module **216** utilizes a Support Vector Machine (SVM) to classify every segment represented in the

context-based feature set **224** as commercial or non-commercial. To this end, a kernel SVM is used to train at least one SVM model **218** for commercial segments. For a segment  $C_j$ , we denote the SVM classification output as  $\text{Cls}(C_j)$ ;  $\text{Cls}(C_j) \geq 0$  indicates that  $C_j$  is a commercial segment. Although techniques to train SVM classification models are known, for purposes of discussion an overview of learning by kernel SVM follows.

Consider the problem of separating a set of training vectors belonging to two separate classes,  $(x_1; y_1), \dots, (x_l; y_l)$ , where  $x_i \in R^n$  is a feature vector and  $y_i \in \{-1, +1\}$  is a class label, with a separating hyper-plane of equation  $w \cdot x + b = 0$ . Of all the boundaries determined by  $w$  and  $b$ , the one that maximizes the margin will generalize better than other possible separating hyper-planes.

A canonical hyper-plane [10] has the constraint for parameters  $w$  and  $b$ :  $\min x_i y_i [(w \cdot x_i) + b] = 1$ . A separating hyper-plane in canonical form must satisfy the following constraints,  $y_i [(w \cdot x_i) + b] \geq 1, i=1, \dots, l$ . The margin is

$$2 / \|w\|$$

according to its definition. Hence the hyper-plane that optimally separates the data is the one that minimizes

$$\phi(w) = \frac{1}{2} \|w\|^2.$$

The solution to the optimization problem is given by the saddle point of the Lagrange functional,

$$L(w, b, \alpha) = \frac{1}{2} \|w\|^2 - \sum_{i=1}^l \alpha_i \{y_i [(w \cdot x_i) + b] - 1\} \quad (11)$$

with Lagrange multipliers  $\alpha_i$ . The solution is given by,

$$\bar{w} = \sum_{i=1}^l \bar{\alpha}_i y_i x_i, \quad (12)$$

$$\bar{b} = -\frac{1}{2} \bar{w} \cdot [x_r + x_s]$$

wherein  $x_r$  and  $x_s$  are support vectors which belong to class +1 and -1, respectively.

In linearly non-separable but nonlinearly separable case, the SVM replaces the inner product  $x \cdot y$  by a kernel function  $K(x; y)$ , and then constructs an optimal separating hyper-plane in the mapped space. According to the Mercer theorem [10], the kernel function implicitly maps the input vectors into a high dimensional feature space. This provides a way to address the difficulties of dimensionality [10].

Possible choices of kernel functions include: (a) Polynomial  $K(x, y) = (x \cdot y + 1)^d$ , where the parameter  $d$  is the degree of the polynomial; (b) Gaussian Radial Basis (GRB) Function:

$$K(x, y) = \exp\left(-\frac{\|x - y\|^2}{2\sigma^2}\right),$$

## 11

where the parameter  $\sigma$  is the width of the Gaussian function; (c) Multi-Layer perception function:  $K(x,y)=\tanh(\kappa(x-y)-\mu)$ , where the  $\kappa$  and  $\mu$  are the scale and offset parameters. In our method, we use the GRB kernel, because it was empirically observed to perform better than other two.

For a given kernel function, the classifier is given by the following equation:

$$f(x) = \text{sgn}\left(\sum_{i=1}^l \bar{a}_i y_i K(x_i, x) + \bar{b}\right) \quad (13)$$

Results of the described SVM classification are represented in FIG. 2 as segment/scene classifications **226**.

#### Post-processing Operations

LBCCD module **216** of FIG. 2 utilizes a number of post-processing operations to increase the accuracy of the SVM-based classification results described above. In particular, scene-grouping techniques are utilized to further refine the LBCCD algorithms. Scene-grouping techniques are applied, at least in part, on an observation that there is typically a substantial similarity in such visual and audio features as color and audio type within commercial blocks, as well as within non-commercial blocks. In light of this, the LBCCD module combines segments into scenes by the method proposed in reference [11], wherein each scene includes all commercial segments or all non-commercial segments. Consecutive commercial scenes and consecutive non-commercial scenes are merged to form a series of commercial blocks and non-commercial blocks.

For instance, let

Shot= $\{C_0, C_1, \hat{\cdot}, C_{N-1}\}$ , N: number of all shots,

Scene= $\{S_0, S_1, \hat{\cdot}, S_{M-1}\}$ , M: number of all scenes, and

$$S_k = \{C_0^k, C_1^k, \Lambda, C_{N_k-1}^k\}, \sum_{k=0}^{M-1} N_k = N, \quad (14)$$

represent all segments of program data **104** (e.g., a TV program) and the scene-grouping results. This scene-grouping algorithm incorporates techniques described in [11], which is hereby incorporated by reference. In particular, refinement of commercial detection by scene-grouping can then be described as scene classification and merging, wherein each scene is classified based on the following rule:

$$Cls(S_k) = \text{sign}\left(\sum_{j=0}^{N_k-1} \text{sign}(Cls(C_j^k))\right) \quad (15)$$

The variable  $\text{sign}(x)$  is a sign function which returns one (1) when  $x \geq 0$  and negative-one (-1) when  $x < 0$ . This rule indicates that if the number of commercial segments in  $S_k$  is not less than half of  $N_k$ , this scene is classified as a commercial scene; otherwise, the scene is classified as a non-commercial scene.

At this point, a number of initial commercial and non-commercial blocks have been so classified. For purposes of discussion, these initial results are represented as an intermediate form of post processing results **228** of FIG. 2, or as other data **230**. Also, for purposes of discussion, these initial blocks

## 12

are still referred to as commercial scenes and non-commercial scenes. To provide further robustness to these scene-grouping results, the LBCCD module **216** of FIG. 2 evaluates the content of the scene groups as a function one or more configurable threshold values. Based on these evaluations, the LBCCD determines whether the initial scene grouping results (and possibly subsequent iterative scene grouping results) should be further refined to better differentiate commercial from non-commercial content.

In particular, and in this implementation, four (4) configurable thresholds are employed by the LBCCD module **216** (FIG. 2) to remove/reconfigure relatively short scenes, double check long commercial scenes for embedded non-commercial content, detect long commercial portions of a non-commercial scene, and refine the boundaries of commercial and non-commercial segments. Application of these criteria may result in scene splitting operations, whereupon for each split operation, properties for the affected scenes are recalculated using equation (15), followed by a merging operation as discussed above.

With respect to criteria to remove short scenes, a scene is considered to be too short if it is smaller than a configurable threshold  $T_1$ . If a scene meets this criterion, the scene is merged into the shorter scene of its two neighbor scenes.

With respect to double checking long commercial scenes, a commercial is not typically very long with respect to the amount of time that it is presented to an audience. Thus, if a commercial scene is longer than a configurable threshold  $T_2$ , the scene is evaluated to determine if it may include one or more non-commercial portion(s). To this end, the LBCCD module **216** determines whether an improperly merged scene or segment (e.g., shot) grouping lies in a boundary between two segments  $C_i$  and  $C_{i+1}$ , according to the following:

$$Cls(C_i) \cdot Cls(C_{i+1}) < 0 \quad (16)$$

$$|Cls(C_i) - Cls(C_{i+1})| > T_2 \quad (17).$$

If these two constraints are satisfied at the same time, the LBCCD module splits the long commercial scene between  $C_i$  and  $C_{i+1}$ . These split scenes are reclassified according to equation (15), followed by a merging operation as discussed above.

With respect to separating a commercial part from a long non-commercial scene, it has been observed that there may be one or more consecutive commercial segments in a long non-commercial scene. To separate any commercial segments from the non-commercial scene in such a situation, the scene is split at the beginning and end of this commercial part, if the number of the consecutive commercial segments is larger than a threshold  $T_c$ . Consecutive commercial segments in a long non-commercial scene are detected by counting the number of consecutive segments that are classified as commercial segments by the aforementioned SVM classification approach. If the number is greater than a configurable threshold, this set of the consecutive segments are regarded as consecutive commercial segments in this long non-commercial scene.

With respect to refining scene boundaries, it is a user-preference rule. If the user wants to keep all non-commercial (commercial) segments, several segments in the beginning and end of each commercial (non-commercial) scene are checked. If a segment  $C_j$  of this kind is too long (short) and

$$Cls(C_j)$$

is smaller (bigger) than a configurable threshold  $T_3$ , the segment is merged/transferred to its closest corresponding non-commercial (commercial) scene.

The client **200** has been described with respect to architectural aspects of a set-top box but may have also been described as a different computing device such as a digital video recorder, a general purpose computing device such as a PC, and so on. Moreover, the client **200** may further include other components, which are not shown for simplicity purposes. For instance, the client may be equipped with hardware and/or software to present a graphical user interface to a viewer, by which the viewer can navigate an electronic program guide (EPG), or (if enabled) to access various Internet system network services, browse the Web, or send email. Other possible components might include a network connection (e.g., modem, ISDN modem, etc.) to provide connection to a network, an IR interface, display, power resources, etc. A remote control may further be provided to allow the user to control the client.

#### Exemplary Procedure

FIG. **6** shows an exemplary procedure for implementing learning-based automatic commercial content detection. For purposes of discussion, the operations of the procedure are described in reference to various program module and data components of FIGS. **1** and **2**. At block **602**, a kernel SVM is used to train one or more SVM classification model(s) for classifying commercial segments. These one or more trained models are represented as trained SVM models **218** (FIG. **2**). At block **604**, the learning-based (LB) commercial detection module **216** (FIG. **2**) segments program data **104** (FIGS. **1** and **2**) into multiple segments. Such segments are represented as extracted segments **220** (FIG. **2**). At block **606**, the LBCCD module analyzes segment-based and time-based visual features as well as audio features of the extracted segments. This analysis generates the visual and audio feature data **222** (FIG. **2**). At block **608**, the results of visual and audio feature analysis are further refined by determining contextual aspects of the segments. Such contextual aspects are represented as context-based feature data **224** (FIG. **2**).

At block **610**, the LBCCD module **216** (FIG. **2**) applies the trained SVM classification models (see, block **602**) to the extracted segments **220** (FIG. **2**) in view of the visual and audio feature data **222** and the context-based feature data **224**. This operation results in each of the extracted segments being classified as commercial or non-commercial. These results are represented as segment/scene classifications **226** (FIG. **2**). At block **612**, the LBCCD module performs a number of post-processing operations to further characterize the segment/scene classifications as being commercial or non-commercial in content. Such post processing operations include, for example, scene-grouping, merging, and application of multiple heuristic criteria to further refine the SVM classifications.

In this manner, the LBCCD module **216** (FIG. **2**) identifies which portions of program data **104** (FIGS. **1** and **2**) consist of commercial/advertising as compared to non-commercial/general programming content. Segregated commercial and/or non-commercial blocks resulting from these post-processing operations are represented as post-processing results **228** (FIG. **2**).

FIG. **7** shows another exemplary procedure for implementing learning-based automatic commercial content detection. For purposes of discussion, the operations of the procedure are described in reference to various program module and data components of FIGS. **1** and **2**. In one implementation, the operations of procedure **700** are implemented by respective components of flash memory **212** of FIG. **2**. Operations of block **702** use a support vector machine to train a classification model to detect commercial segments. Operations of

block **704** analyze program data using a classification model and one or more respective single-side left and/or right neighborhoods to classify segments of the program data as commercial or non-commercial. Operations at block **706** post-process the classified segments into commercial content blocks and non-commercial content box. For example, multiple segments classified as commercial are merged into a commercial content block. Analogously, multiple segments classified as non-commercial can be merged into a non-commercial content block. In one implementation, there are multiple such types of blocks. Operations at block **708** re-analyze the one or more commercial content blocks and non-commercial content blocks to determine whether encapsulated segments of the program data should be reclassified and merged with a different commercial or non-commercial content block. Operations of block **710** indicate to a user or computer-program which portions of program data are commercial and/or non-commercial based on the contents of the commercial and non-commercial content blocks.

#### CONCLUSION

The described systems and methods provide for learning-based automatic commercial content detection. Although the systems and methods have been described in language specific to structural features and methodological operations, the subject matter as defined in the appended claims are not necessarily limited to the specific features or operations described. Rather, the specific features and operations are disclosed as exemplary forms of implementing the claimed subject matter.

#### APPENDIX-REFERENCES

- [1] R. Lienhart, et al. "On the Detection and Recognition of Television Commercials," Proc of IEEE Conf on Multimedia Computing and Systems, Ottawa, Canada, pp. 609-516, June 1997.
- [2] D. Sadlier, et al, "Automatic TV Advertisement Detection from MPEG Bitstream," Intl Conf on Enterprise Information Systems, Setubal, Portugal, 7-10 Jul. 2001.
- [3] T. Hargrove, "Logo Detection in Digital Video," <http://toonarchive.com/logo-detection/>, March 2001.
- [4] R. Wetzel, et al, "NOMAD," <http://www.fatalfx.com/nomad/>, 1998.
- [5] J. M. Sánchez, X. Binefa. "AudiCom: a Video Analysis System for Auditing Commercial Broadcasts," Proc of ICMCS'99, vol. 2, pp. 272-276, Firenze, Italy, June 1999.
- [6] R. Zabih, J. Miller, K. Mai, "A Feature-Based Algorithm for Detecting and Classifying Scene Breaks," Proc of ACM Multimedia 95, San Francisco, Calif., pp. 189-200, November 1995.
- [7] L. Lu, H. J. Zhang, H. Jiang, "Audio Content Analysis for Video Structure Extraction," Submitted to IEEE Trans on SAP.
- [8] L. Lu, H. Jiang, H. J. Zhang. "A Robust Audio Classification and Segmentation Method," 9th ACM Multimedia, pp. 203-211, 2001.
- [9] L. Lu, Stan Li, H, J. Zhang, "Content-based Audio Segmentation Using Support Vector Machines," Proc of ICME 2001, pp. 956-959, Tokyo, Japan, 2001
- [10] V. N. Vapnik, "Statistical Learning Theory", John Wiley & Sons, New York, 1998.
- [11] X. Y. Lu, Y. F. Ma, H. J. Zhang, L. D. Wu, "A New Approach of Semantic Video Segmentation," Submitted to ICME2002, Lausanne, Switzerland, August 2002.



The invention claimed is:

1. A computer-implemented method for learning-based automatic commercial content detection, the method comprising:

dividing program data into multiple segments; 5  
analyzing the segments to determine visual, audio, and context-based feature sets that differentiate commercial content from non-commercial content; wherein the context-based features are a function of one or more single-side left and/or right neighborhoods of segments of the multiple segments; and

calculating context-based feature sets from segment-based visual features as an average value of visual features of  $S^k$ ,  $S^k$  representing a set of all segments of the multiple segments that are partially or totally included in the single-side left and/or right neighborhoods such that  $S^k = \{C_j^k: 0 \leq j < M^k\} = \{C_i: C_i \cap N^k \neq \Phi\}$ ,  $M^k$  being a number of segments in  $S^k$ , and

wherein  $N^k$  represents  $2n+1$  neighborhoods,  $n$  represents a number of neighborhoods left and/or right of a current segment  $C_i$ ,  $S^k$  is a set of segments that are partially or totally included in  $N^k$ ,  $C_j^k$  represents a  $j$ -th element of  $S^k$ ,  $M^k$  represents a total number of elements in  $S^k$ , and  $\Phi$  represents an empty set.

2. A computer-implemented method of claim 1, wherein the visual feature set comprises segment-based features and time-based features.

3. A computer-implemented method of claim 1, wherein the single-side left and/or right neighborhoods comprise a current segment having  $n$  neighborhoods left of the current segment and/or  $n$  neighborhoods right of the current segment.

4. A computer-implemented method of claim 1, wherein the method further comprise:

classifying each of the multiple segments as being commercial or non-commercial content as a function of a trained commercial content classification model, and the visual, audio, and context-based feature sets.

5. A computer-implemented method of claim 1, wherein the method further comprise:

classifying each of the multiple segments as being commercial or non-commercial content as a function of a trained commercial content classification model, and the visual, audio, and context-based feature sets; and

post-processing classified commercial and/or non-commercial segments to generate commercial and/or non-commercial blocks, each block comprising one or more segments of the multiple segments.

6. A tangible computer-readable data storage medium for learning-based automatic commercial content detection, the

computer-readable medium comprising computer-program executable instructions executable by a processor for:

dividing program data into multiple segments;  
analyzing the segments to determine visual, audio, and context-based feature sets that differentiate commercial content from non-commercial content, wherein the context-based features are a function of one or more single-side left and/or right neighborhoods of segments of the multiple segments; and

calculating context-based feature sets from segment-based visual features as an average value of visual features of  $S^k$ ,  $S^k$  representing a set of all segments of the multiple segments that are partially or totally included in the single-side left and/or right neighborhoods such that  $S^k = \{C_j^k: 0 \leq j < M^k\} = \{C_i: C_i \cap N^k \neq \Phi\}$ ,  $M^k$  being a number of segments in  $S^k$ , and

wherein  $N^k$  represents  $2n+1$  neighborhoods,  $n$  represents a number of neighborhoods left and/or right of a current segment  $C_i$ ,  $S^k$  is a set of segments that are partially or totally included in  $N^k$ ,  $C_j^k$  represents a  $j$ -th element of  $S^k$ ,  $M^k$  represents a total number of elements in  $S^k$ , and  $\Phi$  represents an empty set.

7. The tangible computer-readable data storage medium of claim 6, wherein the visual feature set comprises segment-based features and time-based features.

8. The tangible computer-readable data storage medium of claim 6, wherein the single-side left and/or right neighborhoods comprise a current segment, the current segment having  $n$  neighborhoods left of the current segment and/or  $n$  neighborhoods right of the current segment.

9. The tangible computer-readable data storage medium of claim 6, wherein the computer-program instructions further comprise instructions for classifying each of the multiple segments as being commercial or non-commercial content as a function of a trained commercial content classification model, and the visual, audio, and context-based feature sets.

10. The tangible computer-readable data storage medium of claim 6, wherein the computer-program instructions further comprise instructions for:

classifying each of the multiple segments as being commercial or non-commercial content as a function of a trained commercial content classification model, and the visual, audio, and context-based feature sets; and

post-processing classified commercial and/or non-commercial segments to generate commercial and/or non-commercial blocks, each block comprising one or more segments of the multiple segments.

\* \* \* \* \*